Robust Intrinsic and Extrinsic Calibration of RGB-D Cameras

Filippo Basso, Emanuele Menegatti and Alberto Pretto.

arXiv:1701.05748v1 [cs.CV] 20 Jan 2017

***Abstract*—Color-depth cameras (RGB-D cameras) have become the primary sensors in most robotics systems, from service robotics to industrial robotics applications. Typical consumer-grade RGB-D cameras are provided with a coarse intrinsic and extrinsic calibration that generally does not meet the accuracy requirements needed by many robotics applications (e.g., high accuracy 3D environment reconstruction and mapping, high precision object recognition and localization, ...)**

In this paper, we propose a human-friendly, reliable and accurate calibration framework that enables to easily estimate both the intrinsic and extrinsic parameters of a general color-depth sensor couple. Our approach is based on a novel, two components, measurement error model that unifies the error sources of different RGB-D pairs based on different technologies, such as structured-light 3D cameras and time-of-flight cameras. The proposed correction model is implemented using two different parametric undistortion maps that provide the calibrated readings by means of linear combinations of control functions.

Our method provides some important advantages compared to other state-of-the-art systems: it is general (i.e., well suited for different types of sensors), it is based on an easy and stable calibration protocol, it provides a greater calibration accuracy, and it has been implemented within the ROS robotics framework. We report detailed and comprehensive experimental validations and performance comparisons to support our statements.

I. Introduction

The availability of affordable depth sensors in conjunction with common RGB cameras, often embedded in the same device (called RGB-D cameras), has provided mobile robots with a complete and instantaneous representation of both the appearance and the 3D structure of the surrounding environment. Many robotic tasks highly benefit from using such sensors, among others SLAM and navigation [1], [2], tracking [3], [4], object recognition and localization [5] and many others. While color information is typically provided by RGB cameras, there are many technologies able to provide depth information, e.g. time-of-flight (ToF) cameras, laser range scanners and sensors based on structured-light (SL). Even if there are many devices able to provide both color and depth data (RGB-D pairs like the Microsoft Kinect 1 and Kinect 2, the Asus Xtion, the Intel RealSense, etc.), as far as we know, there are no integrated, CMOS-like, imaging sensors able to provide both color and depth information yet. It is clear that, in order to obtain a reliable and accurate scene representation, not only the intrinsic parameters of each sensor should be precisely calibrated, but also the extrinsic parameters relating the two sensors should be precisely known. RGB-D devices are often factory calibrated, with the calibration parameter set stored inside a non-volatile memory. Unfortunately, the quality of such calibrations is only adequate for gaming purposes. For instance, with a default setup, the acquired point clouds can suffer from a non accurate association between depth and RGB data (e.g., see Fig. 26), due to a non perfect alignment between the camera and the depth sensor. Moreover, depth images can suffer from an irregular geometric distortion and a systematic bias in the measurements. A proper calibration method for robust robotics applications is needed.

In this paper, we propose a novel, two-steps calibration method that employs a simple data collection procedure that only needs a minimally structured environment and that does not require any parameters tuning or a great interaction with the calibration software. Moreover, even if the principal targets of the method are the sensors based on structured-light, it is thought to be used also with any RGB camera-depth sensor couple. Given a calibrated RGB camera and an uncalibrated depth sensor, the proposed method automatically infers two correction maps for the depth sensor and, as a "side effect", the rigid body transformation that relates the two sensor reference frames. For the depth sensor, we employ a two components error model that includes a pixel-based distortion error along with a systematic error. The calibrated measurements are obtained employing two calibration maps, each one composed by a matrix of functions of the depth. We propose to represent the first map, used to correct the distortion error, by means of a set of functions, iteratively fitted to the acquired data during a first calibration stage. Then, we include the systematic error and the transformation between the sensors in a second stage of the calibration: at this point, we exploit the plane-to-plane constraints between color and depth data to align the two sensors and to infer the systematic error inside a non-linear optimization framework.

The main contributions of this paper are the following:

- A general and experimentally supported measurements error model, that well describes the error of different depth sensor types in an unified way.
- A spatial/parametric undistortion map that models in a compact and efficient way the distortion effect of structured-light depth sensors.
- A novel optimization framework that aims to estimate the camera-depth sensor rigid displacement along with the parametric model that describes the systematic error on the depth measurements.
- An open source implementation of the proposed method,

Basso is with IT+Robotics. Email: filippo.basso@it-robotics.it. Basso and Menegatti are with the Department of Information Engineering, University of Padova, Italy. Email: emg@dei.unipd.it. Pretto is with the Department of Computer, Control, and Management Engineering "Antonio Ruberti", Sapienza University of Rome, Italy. Email: pretto@dis.uniroma1.it

integrated inside the ROS (Robot Operating System) framework [6]. The code along with a tutorial of the calibration process is available at the following website: http://iaslab-unipd.github.io/rgbd_calibration[1].

An exhaustive set of tests aimed at proving the soundness of the proposed method is reported. We also compare our method with other state-of-the-art systems, using the original implementations provided by their authors: our method appears to be more stable and able to provide the most accurate results. We finally report some experiments of an RGB-D visual odometry system applied to a mobile robot, where we show that the accuracy in the ego-motion estimation highly benefits from using RGB-D data calibrated with our system.

This paper is structured as follows. Related work is reviewed in Sect. II. In Sect. III the error on the depth values provided by the sensor is analyzed and discussed. Sect. IV gives a quick overview of the calibration procedure. The first calibration step is detailed in Sect. V, while Sect. VI describes the second calibration step. The results of the calibration procedure and performance comparisons are reported in Sect. VII. Finally, some conclusions are reported in Sect. VIII.

A. Basic Notations

We use non-bold characters x to represent scalars, bold lower case letters $\mathbf{x}$ to represent vectors, with no distinction between cartesian coordinates and homogeneous coordinates. The coordinates of a point $\mathbf{x}$ with respect to the coordinate frame $\mathcal{A}$ are denoted by ${}^{\mathcal{A}}\mathbf{x}$; ${}^{\mathcal{B}}_{\mathcal{A}}\mathbf{T}$ denotes the homogeneous transformation matrix[2] from the reference frame $\mathcal{A}$ to the frame $\mathcal{B}$, such that ${}^{\mathcal{B}}\mathbf{x} = {}^{\mathcal{B}}_{\mathcal{A}}\mathbf{T}{}^{\mathcal{A}}\mathbf{x}$.
An RGB camera is denoted by $\mathcal{C}$ and it provides an RGB image $\mathcal{I}_{\mathcal{C}}$; a depth sensor is denoted by $\mathcal{D}$ and it provides a depth image $\mathcal{I}_{\mathcal{D}}$. From an RGB image of a scene that contains a checkerboard it is possible to extract the checkerboard corners ${}^{\mathcal{I}_C}B$, where the superscript $\mathcal{I}_{\mathcal{C}}$ explicits the fact that the corners are expressed in 2D pixel coordinates, i.e., ${}^{\mathcal{I}_C}B = \{(u,v)_1, \ldots, (u,v)_n\}$. From a depth image $\mathcal{I}_{\mathcal{D}}$ it is possible to generate a point cloud ${}^{\mathcal{D}}C$, where the superscript $\mathcal{D}$ explicits the fact that the point cloud is expressed with respect to the coordinate frame $\mathcal{D}$, i.e., ${}^{\mathcal{D}}C = \{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$.

II. Related Work

In the recent years, a great number of calibration methods for RGB-D pairs have been proposed. Early methods were intended to estimate only the extrinsic parameteres of a camera and a laser range finder. Mei and Rives [7] addressed the problem of finding the relative pose between a 2D laser range finder and a catadioptric camera by using a planar calibration pattern perceived from both sensors. Scaramuzza *et al.* [8] propose to map 3D range information collected with a 3D tilting laser range finder into a 2D map (called Bearing Angle image) that highlights salient points of a scene thus the user can manually associate points between the two sensors in an easy way. The final extrinsic calibration is then obtained using a Perspective-n-Point (PnP) algorithm followed by a non-linear refinement step.

The availability in the market of affordable depth sensors (structured-light 3D sensors as the Microsoft Kinect and time-of-flight cameras) has greatly increased the interest in depth sensor systems in both the robotics and computer vision communities. Kim *et al.* [9] presented an analysis of the measurement errors of a MESA SwissRanger time-of-flight camera, highlighting the presence of two components: a random noise and a systematic depth error, consistent over time. Using this depth error model, they proposed a three-stages calibration procedure that aims to estimate both the intrinsic and extrinsic parameters of an array of sensors composed by both ToF and video cameras. Jung *et al.* [10] proposed to calibrate extrinsics and intrinsics of a color camera and a ToF camera pair by using a pattern with 4cm-diameter holes that can be simultaneously localized by both sensors. Unfortunately, the error models used in both these methods are conceived for a SwissRanger-like sensor, and can't be easily adapted to depth sensors based on different technologies.

Early works on the Kinect sensor calibration [11], [12], [13], [14] were based on standard RGB camera calibration techniques. They used a checkerboard pattern to calibrate both the RGB and the IR cameras, often blocking the IR projector and illuminating the target with a halogen lamp in order to better highlight the checkerboard corners in the IR camera. Smisek *et al.* [13] showed that Kinect depth sensors are affected by a sort of radially symmetric distortions. To correct such distortion, they estimated a z-correction image built as the pixel-wise mean of the residuals of the plane fitted to the depth data. The z-correction image was subtracted from the z coordinate of the cloud points to obtain the actual depth values. Unfortunately it is currently well known that in the general case this distortion depends on the depth, i.e. becomes much stronger for increasing depths (e.g., see Sect. III). Also Zhang and Zhang [15] and Mikhelson *et al.* [16] exploited a checkerboard pattern but, differently from the previous approaches, they didn't search for checkerboard corners in the IR image. Zhang *et al.* [15] used the fact that points on the checkerboard should be co-planar, and the plane equation can be estimated with the RGB camera. They also proposed to model the depth error of the Kinect sensor by treating the depth value d as a linear function of the real one d^*, that is $d = \mu d^* + \eta$. A limitation of this method comes from the fact that it requires to have a good initialization of the unknown parameters. Mikhelson *et al.* [16] extracted the structural corners from the point cloud derived from the depth images, in order to locate the checkerboard position also in the depth image plane. This approach assumes that the depth image is not affected by any distortion: as mentioned before, this assumption does not always apply for Kinect-like sensors. Moreover, from our experience, extracting structural corners from a point cloud is an operation that sometimes provides poor results.

Herrera *et al.* [17], [18] described a calibration method that exploits a checkerboard pattern attached to a big plane to

[1]The copy and paste function may not work properly with this url due to the underscore symbol.

[2]Here we implicitly assume that the points are expressed using homogeneous coordinates.

calibrate two color cameras and a Kinect-like depth sensor. This method works directly on the raw data provided by the depth sensor (instead of on the metric data) and, alongside the camera-depth sensor relative displacement, it estimates a disparity distortion correction map that depends on the observed disparity. They estimated a coefficient for each pixel $D(u, v) \in \mathbb{R}$ and two global coefficients $\alpha_0, \alpha_1 \in \mathbb{R}$ such that the actual depth value d^* can be computed as

$$d^* = d + D(u, v) \cdot \exp(\alpha_0 - \alpha_1 \cdot d) \ .$$

Moreover, they used the 4 corners of the checkerboard plane as the initial guess of the relative displacement between the cameras and the depth sensor. For short distances their approach seems to obtain good results, as reported also in [19], [20].

An improvement over the work of Herrera *et al.* is the one presented by Raposo *et al.* [21]. They proposed several modifications to the estimation pipeline that allow their algorithm to achieve a calibration accuracy similar to [18] while using less than $1/6$ of the input frames and running in $1/30$ of the time.

Canessa *et al.* [22], instead, proposed to model the depth error by means of a second degree polynomial for each pixel. In their work, authors first estimated the pose of a "virtual" depth sensor with respect to the RGB camera using an incandescent lamp bulb to light the checkerboard and make the depth map saturate in correspondence of the white squares. Then, they positioned a Kinect in front of a plane with a checkerboard attached and acquired a set of images from 0.6 m to 2 m. Finally, they fitted a second degree polynomial to the sample set of every pixel. Actually, the need of an incandescent lamp makes this system quite cumbersome.

Teichman *et al.* [23] proposed a completely different calibration approach for Kinect-like devices: the undistortion map is estimated by means of a SLAM framework, in an unsupervised way. Their algorithm estimates 6 depth multipliers, at $0, 2, \ldots, 10$ meters, and it corrects the depth measurements using a linear interpolation: to the best of our knowledge, this is one of the first approaches that proved to be able to correct depth data at more than 5 meters. The main drawback of their approach is the time it needs to reach a solution: the optimization procedure takes several hours to converge. Moreover, according to [20], it seems not to perform as well as [18] for short distances. Indeed, in [23] close-range measurements (i.e., less than 2 meters) are considered reliable and directly used in the SLAM pipeline to infer the geometry of the scene.

Fiedler et al. [24] investigated the influence of thermal and environmental conditions on the Kinect's data quality. The experiments turned out that variations of the temperature and air draft have a notable influence on Kinect images and range measurements. They derived temperature-related rules to reduce errors in both the calibration and measurement process.

Recently, Di Cicco *et al.* [25] proposed a non-parametric, unsupervised intrinsics calibration method for depth sensors. The best plane fitted to the data is used as a reference, i.e., the average depth is considered reliable. A complete, discretized undistortion map for the depth data is estimated using a machine learning approach. In this approach the extrinsics calibration of a general RGB-D pair is not considered.

Very recently, Staranowicz *et al.* [19] proposed an RGB-D pair calibration algorithm that uses a spherical object moved in front of the camera as a calibration pattern. A limitation of this method is that they focus on the estimation of the rigid displacement between the two cameras only, i.e., the depth error is not corrected.

In our previous work [26], we proposed a spatial/parametric undistortion map that models the distortion effect of Kinect-like depth sensors in a compact way. Our approach uses a calibration pattern similar to the one used in [17]: besides the undistortion map, we estimate the camera-depth sensor alignment along with a parametric model that well describes the systematic error of the depth measurements. Results were very promising, but this method could get stuck in a local minimum if the depth camera intrinsic parameters (focal lengths and central point) were poorly estimated. In this work we improve our previous method by addressing this limitation with a new, more general error model and a new calibration pipeline that refines also the depth camera intrinsic parameters. The calibration protocol employed in our method is inspired by [17] and [23], while our error model has been designed taking inspiration from both the error models presented in [9] and [22].

III. Depth Error Analysis

In this section we introduce our depth error correction model, derived from an experimental analysis performed using two popular structured-light based depth sensors: the Microsoft Kinect and the Asus Xtion Pro Live RGB-D cameras. The choice of such sensors is justified by the fact that they are characterized by the most complex and complete measurements error model.

To analyze the systematic and random errors in the depth measurements, we positioned three SL based sensors (two Microsoft Kinect 1 and an Asus Xtion Pro Live, rigidly mounted on a single support) parallel to a flat wall at increasing distances. We refer to [27] and [28] with regard to the error analysis of a ToF camera: in these papers the authors report a set of comprehensive tests using the Microsoft Kinect 2 ToF camera.

For each position, we collected the sensors readings (i.e., the depth images and the generated point clouds) while measuring the real distances (i.e., the ground truth) using two high precision laser distance meters (e.g., Fig. 11a). Some qualitative results of such analysis are reported in Fig. 1. From our analysis, we noticed that:

1) In the case of the SL sensors, the surfaces defined by the point clouds are not properly planar as they should be, and this "distortion" effect becomes more accentuated for increasing distances (this is the *myopic* property defined in [23]). Moreover, each sensor has a different "distortion pattern". ToF sensors are affected by a negligible distortion effect (usually less than 1 cm), and this error is mainly due to the camera lens distortion.

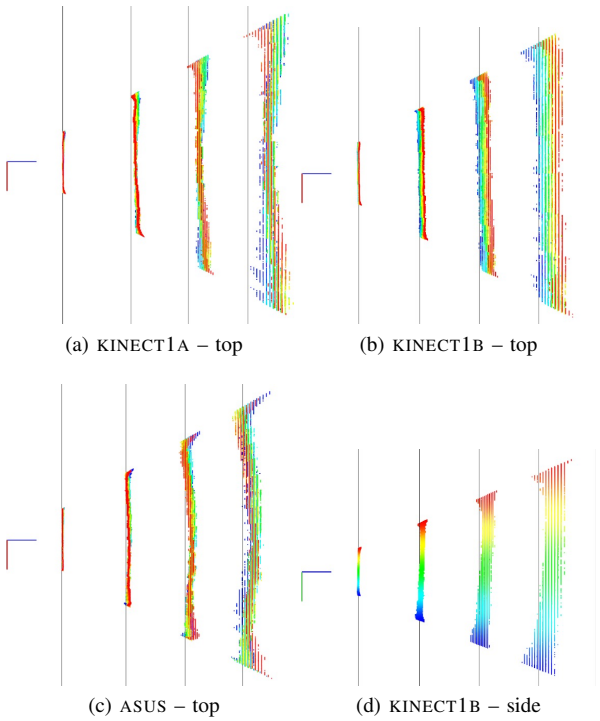

(a) KINECT1A – top (b) KINECT1B – top

(c) ASUS – top (d) KINECT1B – side

Fig. 1. Top and side views of the point clouds generated by some different structured-light depth sensors: two Kinects (KINECT1A, KINECT1B) and an Asus Xtion Pro Live (ASUS). The gray lines represent the ground truth measured with the laser distance meters, points with different y-coordinates are drawn with different colors.

2) The average depth of each point cloud is not generally correct and, sometimes, even the average orientation is wrong.
3) The quantization error for the SL sensors becomes not negligible for increasing distances. ToF sensors are not affected by any quantization error.

The effect of 1) is to produce a local alteration of an object shape, while 2) is a systematic bias in the measurements.

In this work we aim to remove both 1) and 2), while it is not possible to remove the quantization error 3). In the following, we refer to these error components as *distortion error* and *global error*, respectively.

In Fig. 2a we reported the global error of the three SL sensors for increasing distances, i.e. the difference between the average depth of the acquired clouds and the ground truth. It can be noticed that such error is super-linear with respect to the measured depth values. In the case of a ToF sensor, the global error is constant with respect to the measured depth values [28].

We also compared the measured point clouds with the planes that best fit to them (some results are reported in Fig. 2b). In particular, for each incoming point cloud, we computed the Root Mean Square (RMS) error on the (signed) distance between the plane and the points. Also in this case, the error is clearly super-linear.

A. Error Correction Model

To model the effects of the errors introduced by depth sensors, as in [18], [23], [25], [22], we propose to estimate

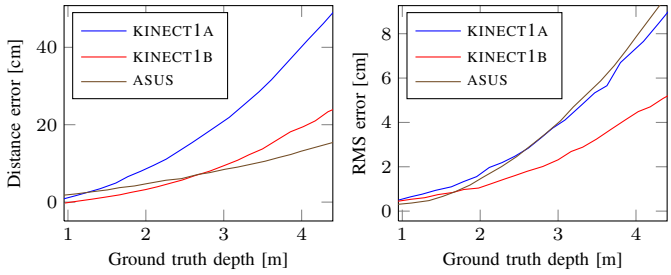

(a) Global error (b) Distortion

Fig. 2. (a) Error on the average distance point estimation, for three different depth sensors. The error is computed, for each cloud, by averaging the depth of all its points and comparing it to the ground truth computed with the two laser distance meters. (b) RMS error caused by the distortion, for three different depth sensors. The error is computed by fitting a plane to the point cloud acquired in front of a flat wall and computing the point-plane distance for all its points.

a depth correction function in a per-pixel basis. That is, given a depth sensor $\mathcal{D}$ that provides a depth image $\mathcal{I}_\mathcal{D}$ of size $H_\mathcal{D} \times W_\mathcal{D}$, a pixel $(u,v)^\mathsf{T}$ and the corresponding depth value $d = \mathcal{I}_\mathcal{D}(u,v)$, the real depth d^* is computed as:

$$d^* = \mathrm{f}_{u,v}(d). \tag{1}$$

Starting from the considerations made in Sect. III, we express each $\mathrm{f}_{u,v}(\cdot)$ in (1) as a composition of two functions: $\mathrm{u}_{u,v}(\cdot)$ that takes into account the local distortion, and $\mathrm{g}_{u,v}(\cdot)$ that makes a global correction of the depth values. That is, the real depth d^* is estimated as:

$$d^* = \mathrm{f}_{u,v}(d) = (\mathrm{g} \circ \mathrm{u})_{u,v}(d), \tag{2}$$

or, alternatively, given the 3D point $\mathbf{x} = (x,y,z)^\mathsf{T}$ associated with the pixel $(u,v)^\mathsf{T}$, the real 3D point $\mathbf{x}^*$ is estimated as:

$$\mathbf{x}^* = \dot{\mathrm{f}}_{u,v}(\mathbf{x}) = (\dot{\mathrm{g}} \circ \dot{\mathrm{u}})_{u,v}(\mathbf{x}),$$

where

$$\dot{\mathrm{f}}_{u,v}(\mathbf{x}) = \mathbf{x} \cdot \frac{\mathrm{f}_{u,v}(d)}{d}.$$

Without loss of generality, we can assume that the depths are already scaled so that $d = \mathcal{I}_\mathcal{D}(u,v) = z$, so in the rest of the paper we will use d and z indifferently.

We define U as the map that associates a pixel $(u,v)^\mathsf{T}$ to an undistortion function $\mathrm{u} : \mathbb{R} \to \mathbb{R}$, that is, $U(u,v) = \mathrm{u}_{u,v}(\cdot)$. In the same way, we define $G(u,v) = \mathrm{g}_{u,v}(\cdot)$ as the map that associates a pixel $(u,v)^\mathsf{T}$ to a function of the depth that correct the global error.

IV. Calibration Approach

As confirmed by the experimental evidence (see Sect. III), the error on the depth measurements is a continuous function. Thus we can assume that given two close 3D points $\mathbf{x}$ and $\mathbf{y}$ along the same direction, i.e $\mathbf{y} = (1+\varepsilon) \cdot \mathbf{x}$ with $\varepsilon \simeq 0$,

$$\mathbf{y}^* = \dot{\mathrm{f}}(\mathbf{y}) = \dot{\mathrm{f}}((1+\varepsilon) \cdot \mathbf{x}) \simeq (1+\varepsilon) \cdot \dot{\mathrm{f}}(\mathbf{x}) = (1+\varepsilon) \cdot \mathbf{x}^*.$$

where $\dot{\mathrm{f}}(\cdot)$ is the error correction function defined in Eq. 2. This means that, if we know how to "correct" a point $\mathbf{x}$ (i.e. we know the correction function parameters for this point),

we can correct close points with a good approximation using the same parameters.

This assumption is the basis of our algorithm to estimate both the undistortion map U and the global error correction map G. Exploiting the fact that both distortion and quantization errors become more severe for increasing distances, we introduce the idea to estimate the distortion error iteratively, starting from short distances and estimating the error for greater distances using as initial guess the current correction parameters.

The proposed calibration framework requires the depth sensor to be coupled with an RGB camera that frames approximately the same scene: the rigid body transformation that relates the two sensors will be estimated while inferring the depth error correction function. It also requires the two sensors to collect data framing a scene that includes a wall with a checkerboard attached on it, at different distances and orientations.

The calibration is performed in two steps: in the first step the algorithm estimates the undistortion map U; only a rough calibration between the camera and the depth sensor is necessary during this step, the checkerboard is used just to have an idea of the wall location. In the second step the global correction map G is computed. Here the checkerboard poses estimated with the (calibrated) RGB camera are used as a ground truth. That is, the undistorted planes estimated in the first step are forced to match the ones defined by the checkerboard. To this end, the real rigid displacement between the RGB camera $\mathcal{C}$ and the depth sensor $\mathcal{D}$ needs to be known. Unfortunately, to estimate the pose of one sensor with respect to the other, a good estimate of their intrinsic parameters is mandatory. One way to satisfy this circular dependency is to estimate the global correction map G and the rigid body displacement ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}$ simultaneously.

At this point a question arises: why the depth error is corrected in two steps?
Actually, the reason is simple. To guarantee the best results, the camera-depth sensor transformation ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}$ and the global correction map G need to be refined together within an optimization framework. Refine inside the same framework a map such as U, with a different function every pixel, is not a feasible solution. On the other hand, the map G, whose scope is to transform planes into planes, is defined by a dozen parameters (see Sect. VI), thus better suited to be efficiently estimated inside an optimization scheme.

A. Pipeline

The algorithm is organized as in Fig. 3. First of all, the checkerboard corners are extracted from all the collected RGB images $\mathcal{I}_{\mathcal{C},k}$ and the organized point clouds[3] are generated from the depth images $\mathcal{I}_{\mathcal{D},k}$, $k = 1, 2, \ldots, M$. The corners ${}^{\mathcal{I}_\mathcal{C}}B_k$ (in pixel coordinates), the point clouds ${}^{\mathcal{D}}C_k$ (in depth sensor reference frame) and the initial camera-depth transformation ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}_0$ are the inputs for the undistortion map estimation

[3]An organized point cloud ${}^{\mathcal{D}}C$ is a point cloud that reflects the depth image structure, i.e. the points are organized into rows and columns as the depth image, $d = \mathcal{I}_\mathcal{D}(u, v) \Rightarrow \mathbf{x} = {}^{\mathcal{D}}C(u, v)$.

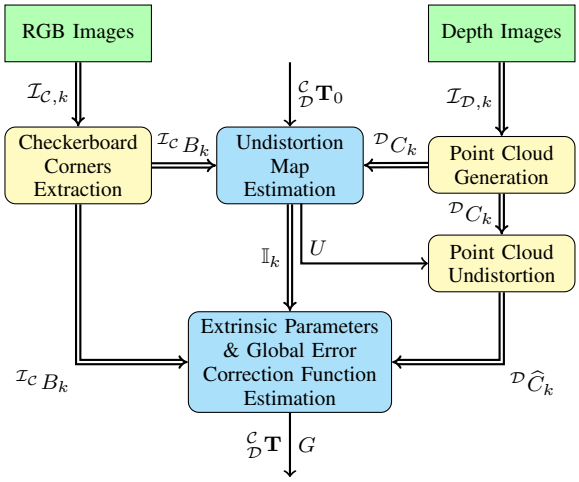

Fig. 3. Calibration algorithm pipeline. Double-lined arrows mean that a set of data is passed from one block to the other.

Input: $\{{}^{\mathcal{D}}C_1, {}^{\mathcal{D}}C_2, \ldots, {}^{\mathcal{D}}C_M\}$ ▷ Point clouds
Input: $\{{}^{\mathcal{I}_\mathcal{C}}B_1, {}^{\mathcal{I}_\mathcal{C}}B_2, \ldots, {}^{\mathcal{I}_\mathcal{C}}B_M\}$ ▷ Checkerboard corners
Input: ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}_0$ ▷ Camera-depth sensors initial transformation
Output: U ▷ Undistortion map

$U \leftarrow \mathtt{1}(\cdot)$ for all $(u, v)^\mathsf{T} \in \mathcal{I}_\mathcal{D}$
$E_U \leftarrow \{\emptyset\}$ for all $(u, v)^\mathsf{T} \in \mathcal{I}_\mathcal{D}$
$\{{}^{\mathcal{D}}C_{s_1}, {}^{\mathcal{D}}C_{s_2}, \ldots, {}^{\mathcal{D}}C_{s_M}\} \leftarrow \textsc{sort}(\{{}^{\mathcal{D}}C_1, {}^{\mathcal{D}}C_2, \ldots, {}^{\mathcal{D}}C_M\})$
for $k \leftarrow 1, 2, \ldots, M$ **do**
  for all $(u, v)^\mathsf{T} \in \mathcal{I}_\mathcal{D}$ **do**
    $\mathtt{u}_{u,v}(\cdot) \leftarrow U(u, v)$
    ${}^{\mathcal{D}}\mathbf{x} \leftarrow {}^{\mathcal{D}}C_{s_k}(u, v)$
    $d \leftarrow \mathcal{I}_{\mathcal{D},s_k}(u, v)$
    ${}^{\mathcal{D}}\widehat{C}_{s_k}(u, v) \leftarrow \dot{\mathtt{u}}({}^{\mathcal{D}}\mathbf{x}) = {}^{\mathcal{D}}\mathbf{x} \cdot \frac{\mathtt{u}_{u,v}(d)}{d}$
  end for
  $\mathbb{I}_{s_k} \leftarrow \textsc{selectWallPoints}({}^{\mathcal{D}}\widehat{C}_{s_k}, {}^{\mathcal{I}_\mathcal{C}}B_{s_k}, {}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}_0)$
  ${}^{\mathcal{D}}\pi_{s_k} \leftarrow \textsc{fitPlane}({}^{\mathcal{D}}C_{s_k}, \mathbb{I}_{s_k})$
  $(U, E_U) \leftarrow \textsc{updateMap}(U, E_U, {}^{\mathcal{D}}C_{s_k}, \mathbb{I}_{s_k}, {}^{\mathcal{D}}\pi_{s_k})$
end for

Fig. 4. Pseudocode of the algorithm developed to estimate the undistortion map U.

module. Once the undistortion map U has been estimated, the point clouds are undistorted (${}^{\mathcal{D}}\widehat{C}_k$) and passed to the module that estimates both the global correction map G and the final camera-depth transformation ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}$.

V. Undistortion Map Estimation

The proposed algorithm (Fig. 4) estimates the undistortion map U taking as input a list of point clouds $\{{}^{\mathcal{D}}C_1, {}^{\mathcal{D}}C_2, \ldots, {}^{\mathcal{D}}C_M\}$ acquired when the depth sensor $\mathcal{D}$ is pointing a planar surface (e.g. a wall) at different distances and orientations. It also requires the positions of the checkerboard corners $\{{}^{\mathcal{I}_\mathcal{C}}B_1, {}^{\mathcal{I}_\mathcal{C}}B_2, \ldots, {}^{\mathcal{I}_\mathcal{C}}B_M\}$, extracted from the images, and a rough estimate of the rigid-body transformation that relates the two sensors ${}^{\mathcal{C}}_{\mathcal{D}}\mathbf{T}_0$. The data structure E_U used in Fig. 4 is a matrix of *sample sets* (one for each sensor pixel) that keeps in memory the samples used to fit the undistortion functions $\mathtt{u}_{u,v}(\cdot)$.

Firstly ($\ell\ell$. 1-2) the undistortion map U is initialized as an $H_\mathcal{D} \times W_\mathcal{D}$ matrix of identity functions $\mathtt{1}(\cdot) : \mathbb{R} \to \mathbb{R}$ while

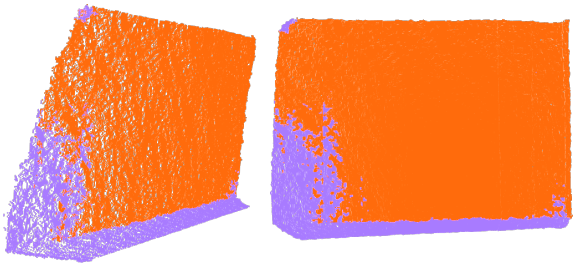

(a) Wall points extracted (in orange) from the original point cloud ${}^{\mathcal{D}}C_{s_k}$.

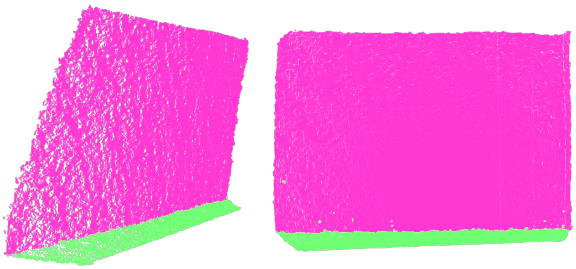

(b) Wall points extracted (in pink) from the undistorted point cloud ${}^{\mathcal{D}}\widehat{C}_{s_k}$.

Fig. 5. Comparison between (a) the wall points extracted from the original point cloud ${}^{\mathcal{D}}C_{s_k}$ and (b) the ones extracted from the undistorted cloud ${}^{\mathcal{D}}\widehat{C}_{s_k}$. As clearly visible, in both cases the floor points are correctly discarded. In the former case, however, the wall segmentation is wrong.

the sample matrix E_U is initialized as an $H_{\mathcal{D}} \times W_{\mathcal{D}}$ matrix of empty sets. Then, the point cloud list $\{{}^{\mathcal{D}}C_1, {}^{\mathcal{D}}C_2, \dots, {}^{\mathcal{D}}C_M\}$ is sorted in ascending order, $\{{}^{\mathcal{D}}C_{s_1}, {}^{\mathcal{D}}C_{s_2}, \dots, {}^{\mathcal{D}}C_{s_M}\}$, according to the distance of the main plane (i.e., the plane with the checkerboard) from the sensor (ℓ. 3), to exploit the continuity described in Sect. IV.

The undistortion map is created iteratively: at each step only one point cloud is analyzed. At step k, for example, the k^{th} cloud in the sorted list, ${}^{\mathcal{D}}C_{s_k}$, is undistorted using the *current* estimation of U ($\ell\ell$. 5-9). Then, the coordinates $\mathbb{I}_{s_k} = \{(u,v)^{\mathsf{T}}_{s_k,1}, (u,v)^{\mathsf{T}}_{s_k,2}, \dots, (u,v)^{\mathsf{T}}_{s_k,n_{s_k}}\}$ of the n_{s_k} points that lie to the main plane are extracted from the undistorted cloud ${}^{\mathcal{D}}\widehat{C}_{s_k}$ and used to fit a plane ${}^{\mathcal{D}}\pi_{s_k}$ to the initial cloud ${}^{\mathcal{D}}C_{s_k}$ ($\ell\ell$. 10-11, see also Sect. V-A). Actually, to increase stability, instead of fitting a plane to the whole original point cloud, only the pixels within a defined ray from the plane center are used, as reported in [25]. Finally (ℓ. 12), the estimated plane ${}^{\mathcal{D}}\pi_{s_k}$ is used to compute the undistortion map U. The procedure ends as soon as the last cloud in the list has been processed.

A. Wall Points Selection

The selection of the wall point-coordinates $\mathbb{I}_{s_k}$ is performed automatically, as opposed to the manual selection of [18]. We take advantage of the RGB camera and the checkerboard to select the right plane and extract the coordinates from the undistorted cloud with a RANSAC-based approach [29], [30]. As the example shown in Fig. 5, the undistorted cloud lets us extract the correct points, where the original one does not. In particular, points near the image corners are likely to be excluded from the inliers when using the original cloud.

1: **function** UPDATEMAP($U, E_U, {}^{\mathcal{D}}C, \mathbb{I}, {}^{\mathcal{D}}\pi$)
2: **for all** $(u,v)^{\mathsf{T}} \in \mathbb{I}$ **do**
3: ${}^{\mathcal{D}}\mathbf{x} \leftarrow {}^{\mathcal{D}}C(u,v)$
4: ${}^{\mathcal{D}}\mathbf{x}_\pi \leftarrow$ LOSPROJECT$({}^{\mathcal{D}}\mathbf{x}, {}^{\mathcal{D}}\pi)$
5: $E_U(u,v) \leftarrow E_U(u,v) \cup \{(z, z_\pi)\}$
6: $U(u,v) \leftarrow$ FITCURVE$(E_U(u,v))$
7: **end for**
8: **return** (U, E_U)
9: **end function**

Fig. 6. Pseudocode of the algorithm for updating the undistortion map U.

B. Map Update

In the map update function (Fig. 6), all the points ${}^{\mathcal{D}}\mathbf{x} = (x,y,z)^{\mathsf{T}}$ of the cloud that lie to the main plane are projected on the previously extracted plane ${}^{\mathcal{D}}\pi_{s_k}$ along their line-of-sight ($\ell\ell$. 3-4, where LOSPROJECT() is a short name for the function LINEOFSIGHTPROJECT()). That is, let $\mathbf{n}^{\mathsf{T}}\mathbf{x} - d_\pi = 0$ be the plane equation, and let $l{}^{\mathcal{D}}\mathbf{x}$, $l \in \mathbb{R}$, be the points along ${}^{\mathcal{D}}\mathbf{x}$-line-of-sight, then the line-of-sight projection of ${}^{\mathcal{D}}\mathbf{x}$ onto ${}^{\mathcal{D}}\pi_{s_k}$, say ${}^{\mathcal{D}}\mathbf{x}_\pi = (x_\pi, y_\pi, z_\pi)^{\mathsf{T}}$, is:

$$\mathbf{x}_\pi = l\mathbf{x} = \frac{d_\pi \mathbf{x}}{\mathbf{n}^{\mathsf{T}}\mathbf{x}} \ .$$

where the superscript $\mathcal{D}$ has been omitted for the sake of simplicity. The pair (z, z_π) is used as a sample for the curve-fitting procedure (ℓ. 5), and the undistortion function $U(u,v)$ is re-estimated by fitting a new curve to the sample set $E_U(u,v)$ (ℓ. 6).

C. Implementation Details

1) Undistortion Map: To decrease the incidence of noise on the map estimation we reduce the number of functions fitted to the data. That is, instead of estimating an undistortion function for each pixel, similarly to [23], we discretize the map into bins. So, let $\chi_U, \psi_U \in \mathbb{N}$ be the bin size in pixels, along the image x- and y-direction, respectively.

Given a pixel $(u,v)^{\mathsf{T}} \in \mathcal{I}_{\mathcal{D}}$, we define $S_U(u,v)$ as the set of 4 pixels *surrounding* $(u,v)^{\mathsf{T}}$ according to the sampling factors χ_U and ψ_U (see Fig. 7). We also define $\mathbb{S}_U \triangleq \{S_U(u,v) : (u,v)^{\mathsf{T}} \in \mathcal{I}_{\mathcal{D}}\}$ as the set of all the *surrounding pixels*.

We estimate the undistortion function $\mathrm{u}_{u,v}(\cdot)$ only for the pixels in $\mathbb{S}_U$. For all the others, instead, this function is computed as a linear combination of the functions computed for the pixels set $\mathbb{S}_U$. That is, given a pixel $(u,v)^{\mathsf{T}}$, following a bilinear interpolation approach, its undistortion function $U(u,v)$ can be computed as:

$$U(u,v) = \sum_{(s,t)\in S_U(u,v)} \mathrm{w}_{\chi_U}(u,s) \cdot \mathrm{w}_{\psi_U}(v,t) \cdot U(s,t)$$

where

$$\mathrm{w}_{\chi_U}(u,s) \triangleq 1 - \frac{|u-s|}{\chi_U}, \quad \mathrm{w}_{\psi_U}(v,t) \triangleq 1 - \frac{|v-t|}{\psi_U} \tag{3}$$

and

$$\sum_{(s,t)\in S_U(u,v)} \mathrm{w}_{\chi_U}(u,s) \cdot \mathrm{w}_{\psi_U}(v,t) = 1 \ .$$

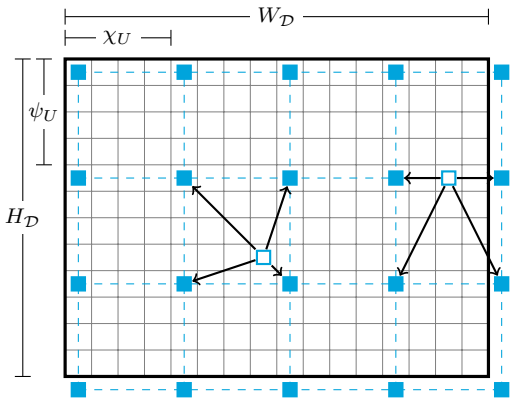

- Pixel $(u,v)^\mathsf{T} \in \mathbb{S}_U$. For this pixel an undistortion function $\mathrm{u}_{u,v}(\cdot)$ has been estimated.
- Pixel $(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}}$. For this pixel the undistortion function $\mathrm{u}_{u,v}(\cdot)$ is a linear combination of the functions of the pixels in $S_U(u,v)$.
- → Connection from a pixel $(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}}$ to one in $S_U(u,v)$.

Fig. 7. Visualization of a discretized undistorted map U. Given two parameters, χ_U and ψ_U, an undistortion function is estimated for all and only the pixels in $\mathbb{S}_U$. For all the others, instead, the function is computed as a linear combination of the neighboring points in $\mathbb{S}_U$.

2) Curve Fitting: As shown in Sect. III (Fig. 2b), the distortion is super-linear, therefore an appropriate correction function must be chosen. Moreover, as previously described, since we are not estimating a function for every pixel, the fitting procedure is not straightforward.

For what concerns the former point, suppose the error is corrected by a second degree polynomial, that is, $U(u,v) = \mathrm{u}_{u,v}(z) = a + bz + cz^2$, for some $a, b, c \in \mathbb{R}$. To estimate the polynomial coefficients we solve a non-linear least squares problem of the form:

$$\arg\min_{a,b,c} \sum_{(z,z_\pi)\in E_U(u,v)} \frac{1}{\sigma^2(z)} \left\| a + bz + cz^2 - z_\pi \right\|^2$$

where $\sigma(z)$ is a normalization term. For the Kinect 1 we choose $\sigma(z)$ to be the quantization error (reported in [13]), i.e.

$$\sigma(z) = -0.00029 + 0.00037 \cdot z + 0.001365 \cdot z^2.$$

For what concerns the latter point, i.e. how to deal with the discretized undistortion map, we slightly modify the sample set generation and the function fitting procedure described in Fig. 6. In the new algorithm (Fig. 8), every pixel $(u,v)^\mathsf{T}$ contributes to the sample set of its four surrounding pixels $S_U(u,v)$ with a weight calculated as in (3). That is, let:

$$S_U^{-1}(s,t) \triangleq \{(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}} : (s,t) \in S_U(u,v)\}$$

be the set of pixels which have (s,t) as one of their surrounding pixels. For each cloud, the temporary sample set $E_w(s,t)$ for a pixel $(s,t) \in \mathbb{S}_U$, is ($\ell\ell$. 2-10):

$$E_w(s,t) \triangleq \bigcup_{(u,v)^\mathsf{T}\in S^{-1}(s,t)} (w, z, z_\pi)$$

```
1:  function UPDATEMAP(U, E_U, ^D C, I, ^D π)
2:      E_w ← {∅} for all (u,v)^T ∈ I_D
3:      for all (u,v)^T ∈ I do
4:          ^D x ← ^D C(u,v)
5:          ^D x_π ← LOSPROJECT(^D x, ^D π)
6:          for all (s,t) ∈ S_U(u,v) do
7:              w ← w_χU(u,s) · w_ψU(v,t)
8:              E_w(s,t) ← E_w(s,t) ∪ {(w, z, z_π)}
9:          end for
10:     end for
11:     for all (s,t) ∈ S_U do
12:         (z̄, z̄_π) ← WEIGHTEDMEAN(E_w(s,t))
13:         E_U(s,t) ← E_U(s,t) ∪ {(z̄, z̄_π)}
14:         U(s,t) ← FITCURVE(E_U(s,t))
15:     end for
16:     return (U, E_U)
17: end function
```

Fig. 8. Pseudocode of the algorithm for updating the undistortion map U taking into account the pixel binning.

where

$$w \triangleq \mathrm{w}_{\chi_U}(u,s) \cdot \mathrm{w}_{\psi_U}(v,t).$$

$E_w(s,t)$ is used to generate the sample set for the aforementioned curve fitting procedure ($\ell\ell$. 11-15). Basically, the pair $(\bar{z}, \bar{z}_\pi)$ is calculated as the weighted arithmetic mean of the values in $E_w(s,t)$, that is:

$$\begin{aligned} W &\triangleq \sum_{(w,z,z_\pi)\in E_w(u,v)} w, \\ \bar{z} &\triangleq \frac{1}{W} \sum_{(w,z,z_\pi)\in E_w(u,v)} w \cdot z, \\ \bar{z}_\pi &\triangleq \frac{1}{W} \sum_{(w,z,z_\pi)\in E_w(u,v)} w \cdot z_\pi, \end{aligned}$$

and added to the sample set $E_U(s,t)$.

VI. Global Correction Map Estimation

Our original solution to deal with the global, systematic error was to have a unique function, say $\mathrm{g}(\cdot)$, to correct the wrong depth measurements after the undistortion phase, i.e. $G(u,v) = \mathrm{g}(\cdot)$, for all $(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}}$. Unfortunately, this solution had one important limitation: in some cases the undistorted clouds were both translated and rotated around a non-predictable axis. For this reason we moved to a correction map G someway similar to the previously described undistortion map U. The actual implementation of such map is described in Sect. VI-C1. Our algorithm takes as input a set of already undistorted point clouds $\{{}^{\mathcal{D}}\widehat{C}_1, {}^{\mathcal{D}}\widehat{C}_2, \ldots, {}^{\mathcal{D}}\widehat{C}_M\}$, the correspondent wall point coordinates $\{\mathbb{I}_1, \mathbb{I}_2, \ldots, \mathbb{I}_M\}$ and the checkerboard corners extracted from the RGB images, $\{{}^{\mathcal{I}_\mathcal{C}}B_1, {}^{\mathcal{I}_\mathcal{C}}B_2, \ldots, {}^{\mathcal{I}_\mathcal{C}}B_M\}$. After an initialization step where a rough estimate of the map G is computed (Sect. VI-A), the map is refined, along with the camera-depth sensor transformation ${}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}$, within a non-linear optimization framework (Sect. VI-B).

Input: $\{{}^{\mathcal{D}}\widehat{C}_1, {}^{\mathcal{D}}\widehat{C}_2, \dots, {}^{\mathcal{D}}\widehat{C}_M\}$ ▷ Undistorted point clouds
Input: $\{\mathbb{I}_1, \mathbb{I}_2, \dots, \mathbb{I}_M\}$ ▷ Wall point coordinates
Input: $\{{}^{\mathcal{C}}B_1, {}^{\mathcal{C}}B_2, \dots, {}^{\mathcal{C}}B_M\}$ ▷ Checkerboard corners
Output: G ▷ Global error correction map
Output: ${}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}$ ▷ Camera-depth sensor transformation matrix

1: **for** $k \leftarrow 1, 2, \dots, M$ **do**
2: ${}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T} \leftarrow \text{SOLVEPNP}(\mathbf{K}_{\mathcal{C}}, \mathbf{dist}_{\mathcal{C}}, {}^{\mathcal{B}}B, {}^{\mathcal{I}_{\mathcal{C}}}B_k)$
3: ${}^{\mathcal{C}}B_k \leftarrow {}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T} \cdot {}^{\mathcal{B}}B$
4: ${}^{\mathcal{C}}\pi_{B_k} \leftarrow \text{FITPLANE}({}^{\mathcal{C}}B_k)$
5: ${}^{\mathcal{D}}\pi_{\widehat{C}_k} \leftarrow \text{FITPLANE}({}^{\mathcal{D}}\widehat{C}_k, \mathbb{I}_k)$
6: **end for**
7: $\mathbf{\Pi}_B \leftarrow ({}^{\mathcal{C}}\pi_{B_1}, {}^{\mathcal{C}}\pi_{B_2}, \dots, {}^{\mathcal{C}}\pi_{B_M})$
8: $\mathbf{\Pi}_{\widehat{C}} \leftarrow ({}^{\mathcal{D}}\pi_{\widehat{C}_1}, {}^{\mathcal{D}}\pi_{\widehat{C}_2}, \dots, {}^{\mathcal{D}}\pi_{\widehat{C}_M})$
9: ${}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T} \leftarrow \text{ESTIMATETRANSFORM}(\mathbf{\Pi}_B, \mathbf{\Pi}_{\widehat{C}})$
10: **for** $k \leftarrow 1, 2, \dots, M$ **do**
11: ${}^{\mathcal{D}}\pi_{B_k} \leftarrow {}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T} \cdot {}^{\mathcal{C}}\pi_{B_k}$
12: $(G, E) \leftarrow \text{UPDATEMAP}(G, E, {}^{\mathcal{D}}\widehat{C}_k, \mathbb{I}_k, {}^{\mathcal{D}}\pi_{B_k})$
13: **end for**

Fig. 9. Pseudocode of the algorithm for the initial estimation of the global correction map G.

A. Initial Estimation

The algorithm used for the initial estimation of the map functions as well as the computation of the rigid transform between the RGB and the depth sensor is reported in Fig. 9. Firstly, the pose of one sensor with respect to the other is estimated, that is, for each color-depth image pair both the plane defined by the checkerboard in the image ${}^{\mathcal{C}}\pi_{B_k}$ (in the RGB camera $\mathcal{C}$ reference frame), and the one extracted from the point cloud ${}^{\mathcal{D}}\pi_{\widehat{C}_k}$ (in the depth sensor $\mathcal{D}$ reference frame) are computed from the given input data ($\ell\ell$. 1-6). The transformation between the checkerboard and the RGB camera (${}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T}$) is estimated given the checkerboard 3D points ${}^{\mathcal{B}}B$, their corresponding image projections ${}^{\mathcal{I}_{\mathcal{C}}}B_k$, the camera matrix $\mathbf{K}_{\mathcal{C}}$ and the distortion coefficients $\mathbf{dist}_{\mathcal{C}}$ using an iterative optimization based on the Levenberg-Marquardt method (OpenCV [31] function *solvePnP*, ℓ. 2, that solves a Perspective-n-Point problem) The checkerboard 3D points are then transformed in the RGB camera frame (ℓ. 3). The equation of the plane framed by the RGB camera is hence computed taking 3 non-collinear corners (ℓ. 4). The equation of the plane in the depth image, instead, is computed using a SVD approach (ℓ. 5). Once all the planes have been computed, the rigid displacement between the two sensors is estimated ($\ell\ell$. 7-9) using the plane-to-plane calibration method described in [32].

The plane equations extracted from the images are then represented w.r.t. the depth sensor reference frames using the transformation matrices just computed. These planes are used as reference locations for the curve fitting procedure ($\ell\ell$. 10-13), as we did with the undistortion map U in Fig. 4.

B. Non-linear Refinement

Once the global correction map G and the camera-depth sensor transformation matrix ${}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}$ have been estimated, we refine them within a non-linear optimization framework. To take into account the error on the checkerboard poses estimation, we follow the bundle-adjustment approach as described in [33]: we also refine *all* the checkerboard poses ${}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T}$, with $k = 1, \dots, M$. Moreover, to take into account the error on the intrinsic parameters of the depth camera $\mathbf{K}_{\mathcal{D}}$, the focal lengths and the principal point are refined too. So, let define ${}^{C}\mathbb{T}_{\mathcal{B}} \triangleq \left\{{}^{\mathcal{C}}_{\mathcal{B}_1}\mathbf{T}, {}^{\mathcal{C}}_{\mathcal{B}_2}\mathbf{T}, \dots, {}^{\mathcal{C}}_{\mathcal{B}_M}\mathbf{T}\right\}$ as the set of checkerboard poses in the camera coordinates, estimated with the *solvePnP* function.
Formally, the results of the non-linear refinement is:

$$\left(G, {}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}, {}^{C}\mathbb{T}_{\mathcal{B}}, \mathbf{K}_{\mathcal{D}}\right) = \underset{G, {}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}, {}^{C}\mathbb{T}_{\mathcal{B}}, \mathbf{K}_{\mathcal{D}}}{\arg\min} \sum_{k=1}^{M} e_{\text{repr}}(k) + e_{\text{pos}}(k) \ ,$$

Here e_{repr} takes into account the reprojection error of the checkerboard corners onto the images and depends on the checkerboard poses only. This error component is defined as:

$$e_{\text{repr}}(k) \triangleq \sum_{(r,c)\in\mathcal{B}} \frac{1}{\sigma^2_{\mathcal{C}}} \cdot \left\| \text{proj}_{\mathcal{C}}\left({}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T} \cdot {}^{\mathcal{B}}B_k(r,c)\right) - {}^{\mathcal{I}_{\mathcal{C}}}B_k(r,c) \right\|^2$$

The summation is performed over all the checkerboard corners, $\text{proj}_{\mathcal{C}}$ is a general projection function that depends on both the camera matrix $\mathbf{K}_{\mathcal{C}}$ and the distortion coefficients $\mathbf{dist}_{\mathcal{C}}$. The residuals are weighted by the inverse of the variance of the corner estimation error $\sigma^2_{\mathcal{C}}$, where $\sigma_{\mathcal{C}} = 0.2$.

e_{pos} represents the error between the planes defined by the checkerboards and the ones defined by the undistorted point cloud

$$e_{\text{pos}}(k) \triangleq \sum_{(u,v)\in\mathbb{I}_k} \frac{1}{|\mathbb{I}_k| \cdot \sigma^2_U(z)} \cdot \cdot \left\| \text{p}_{{}^{\mathcal{D}}\pi_k}\left(\text{g}_{u,v}\left({}^{\mathcal{D}}\widehat{C}_k(u,v)\right)\right) - \text{g}_{u,v}\left({}^{\mathcal{D}}\widehat{C}_k(u,v)\right) \right\|^2 \ .$$

where $\text{p}_{\pi}(\mathbf{x})$ is the function that orthogonally projects a point $\mathbf{x}$ onto plane π. Formally speaking, this error is the distance between the cloud point ${}^{\mathcal{D}}\widehat{C}_k(u,v)$ corrected with the current estimation of G and its line-of-sight projection onto the plane ${}^{\mathcal{D}}\pi_k$ defined by the checkerboard corner set ${}^{\mathcal{D}}B_k$. Such set is computed as

$${}^{\mathcal{D}}B_k = {}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T} \cdot {}^{\mathcal{C}}_{\mathcal{B}_k}\mathbf{T} \cdot {}^{\mathcal{B}}B \ .$$

Finally, each residual is weighted by the inverse of the variance on the depth measurements *after* the undistortion phase $\sigma^2_U(z)$, multiplied by the number of wall points, i.e. $|\mathbb{I}_k|$.

C. Implementation Details

1) Global Correction Map: As mentioned before, the global correction map G is someway similar to the undistortion map U, but more simple. In fact, G needs to transform planes into planes and, recalling that a plane transformation has 3 degrees of freedom, we just need 3 functions to satisfy this requirement. Thus we define G as a discretized map constructed as U with $\chi_G = W_{\mathcal{D}}$ and $\psi_G = H_{\mathcal{D}}$, that is, only 4 pixels contain a correction function $\text{g}_{u,v}(\cdot)$. For what concerns the other pixels, the global correction function is computed as a linear combination of the functions for the 4 boundary pixels. Allowing 4 pixels to control the whole map usually leads to wrong results. Let suppose, for example that 3 of such pixels contain an identity function and the fourth does not. Clearly the resulting surface will not be a plane anymore. For this

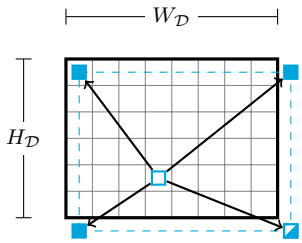

- Pixel $(u,v)^\mathsf{T} \in \mathbb{S}_G$. For this pixel a correction function $\mathrm{g}_{u,v}(\cdot)$ has been estimated.
- Pixel $(u,v)^\mathsf{T} \in \mathbb{S}_G$. To guarantee planarity, for this pixel the correction function $\mathrm{g}_{u,v}(\cdot)$ has been computed starting from the ones of the other 3 pixels in $\mathbb{S}_G$.
- Pixel $(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}}$. For this pixel the correction function $\mathrm{g}_{u,v}(\cdot)$ is a linear combination of the functions of the pixels in $S_G(u,v)$.

→ Connection from a pixel $(u,v)^\mathsf{T} \in \mathcal{I}_{\mathcal{D}}$ to one in $S_G(u,v)$.

Fig. 10. Visualization of a global correction map G. A global correction function is estimated for all and only the pixels in $\mathbb{S}_G$. For all the others pixels, instead, the function is computed as a linear combination of the ones estimated for the pixels in $\mathbb{S}_G$.

reason, only 3 of these pixels are actually computed, the fourth is instead estimated exploiting the following invariant:

$$\frac{\mathrm{g}_{0,0}(d) + \mathrm{g}_{W_{\mathcal{D}},H_{\mathcal{D}}}(d)}{2} = \frac{\mathrm{g}_{W_{\mathcal{D}},0}(d) + \mathrm{g}_{0,H_{\mathcal{D}}}(d)}{2} \quad . \tag{4}$$

Suppose now that $\mathrm{g}_{W_{\mathcal{D}},H_{\mathcal{D}}}(\cdot)$ is the dependent function: $\mathrm{g}_{W_{\mathcal{D}},H_{\mathcal{D}}}(\cdot)$ can be estimated by fitting a function on an adequate set of pairs $(d, \mathrm{g}_{W_{\mathcal{D}},H_{\mathcal{D}}}(d))$, where

$$\mathrm{g}_{W_{\mathcal{D}},H_{\mathcal{D}}}(d) = \mathrm{g}_{W_{\mathcal{D}},0}(d) + \mathrm{g}_{0,H_{\mathcal{D}}}(d) - \mathrm{g}_{0,0}(d)$$

is computed from (4). An example of the presented global correction map G is reported in Fig. 10.

2) Curve Fitting: Since the correction map G is constructed in the same way as the undistortion map U, the considerations on the curve fitting procedure made in Sect. V-C2 are still valid in the global error case.

VII. Experimental Evaluation

The goal of the presented experimental evaluations is to show that our method is able to provide robust and state-of-the-art calibration results for different types of RGB-D sensors. We used four RGB-D sensors: two Microsoft Kinect 1 (called KINECT1A and KINECT1B in the plots), an Asus Xtion Pro Live (called ASUS in the plots), and a Microsoft Kinect 2. The RGB camera of each device has been previously calibrated exploiting a standard calibration tool. Actually, a good RGB camera calibration is an essential requirement of all the tested systems. Each sensor has been mounted, one at a time, on a support that includes two rigidly mounted high precision laser meters, and a high resolution RGB camera (see Fig. 11a).

We attached a checkerboard on a wall, collecting for each device two datasets: a *training set* (i.e., the dataset used to perform the calibration) and a *test set* (i.e., the dataset used to evaluate the calibration accuracy). The training set contains views of the checkerboard from the device camera, the depth sensor and the high resolution camera from different

(a) (b)

Fig. 11. (a) The support used to acquire the data for the experimental evaluations. The two laser meters are located on the left and right of the support to guarantee that the sensor is correctly aligned; (b) The reference hollow cube used as a ground truth in the performance evaluations.

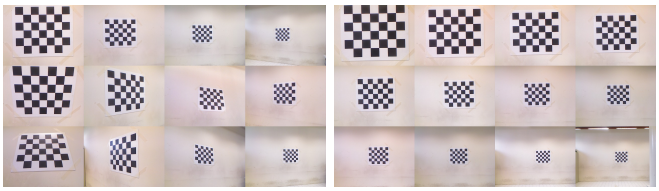

(a) (b)

Fig. 12. (a) Some images of a training set acquired with the sensor KINECT1A. The checkerboard is framed from different locations and orientations from 1 to about 4 meters; (b) Some images of a test set taken with the sensor KINECT1A. The checkerboard is framed from different locations from 1 to about 4.5 meters, paying attention to always have a zero rotation with respect to the wall plane.

locations and orientations (e.g., see Fig. 12a). The test set, instead, has been acquired by positioning the support of Fig. 11a orthogonal to the wall at different distances and measuring such distances with the two laser distance meters (e.g., Fig. 12b).

We independently analyze the results of our undistortion approach (Sect. VII-A), the results of our global correction approach (Sect. VII-B), and the provided camera-depth sensor transformation accuracy (Sect. VII-C). We also compare our method with other state-of-the-art calibration systems (Sect. VII-D), using the original implementations provided by the authors and a reference pattern as a ground truth (Fig. 11b). In almost all tests, our system outperforms the other evaluated systems. We finally report some experiments of an RGB-D visual odometry system applied to a mobile robot, where we show that the accuracy in the ego-motion estimation highly benefits from using RGB-D data calibrated with our system (Sect. VII-E).

A. Undistortion Map

To evaluate the performance of our undistortion approach, we introduce a metric called *planarity error*. For each cloud of the test set, we extract the wall point indices $\mathbb{I}$ from its undistorted version as described in Sect. V-A. We define the planarity error as:

$$e_{\text{plan}} = \sqrt{\frac{1}{|\mathbb{I}|} \sum_{(u,v)\in\mathbb{I}} \left\| \mathbf{n}^\mathsf{T} C(u,v) - d \right\|^2} \quad ,$$

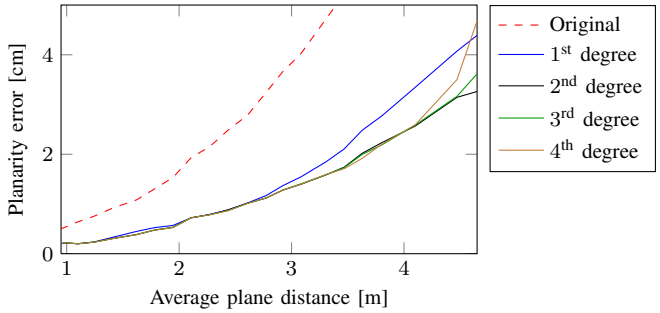

Fig. 13. Planarity error when varying the degree of the undistortion polynomials. From the plot we can see that linear functions are not able to correctly model the distortion introduced by the sensor. On the other hand, quartic functions tend instead to overfit the training data (e.g., in the right part of the plot the error of the quartic functions increases w.r.t. the errors of the quadratic and cubic functions).

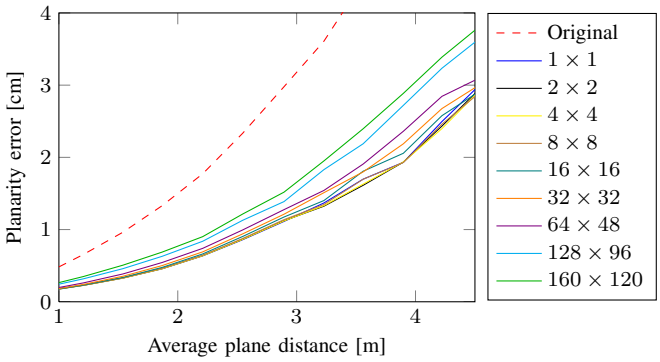

Fig. 14. Planarity error of the wall points when changing the bin size. In the plot it is visible that increasing the bin size also the error increases, but not as much as expected. Actually, only with bin sizes staring from 16×16 the error increase starts being non-negligible.

where C is a generic point cloud of the test set (we consider both the original and the undistorted versions) and π is the plane with equation $\mathbf{n}^\mathsf{T}\mathbf{x} - d = 0$ fitted to the wall points with indices in $\mathbb{I}$.

1) Undistortion Map Functions: In the previous sections, we have always talked about "undistortion functions" without providing many details about the nature of these functions. Actually, analyzing the error on the plane estimation described in Sect. III, we evinced that: (i) for SL sensors, this error is well described by a *quadratic polynomial*; (ii) ToF sensors are affected by an almost negligible distortion effect. The hypothesis (i) is further confirmed also by our experiments. We calibrated each sensor using different types of undistortion functions (linear, quadratic, cubic, ...); for each cloud in the test sets, we computed the planarity error introduced above. The plot in Fig. 13 clearly shows that all the super-linear functions provide better undistortion results when compared to the linear functions. Moreover, higher degree polynomials get similar of even worst results compared with quadratic functions, since they tend to overfit the training data. Therefore, all the tests presented in following sections have been performed using quadratic undistortion functions: in our experience, for SL sensors, this class of functions provides good undistortion results without presenting overfitting problems.

2) Map Discretization: To select the most appropriate bin size values (i.e., χ_U and ψ_U, described in Sect. V-C1), we evaluated the planarity error varying the two parameters. The results are reported in Fig. 14. Differently from what we expected, such parameters do not affect so much the results. Actually, up to a 8×8 pixels size, the planarity error is almost identical. Only with greater sizes, starting from 16×16, the error increases, especially close to the image corners (a qualitative comparison is reported in Fig. 15). Some examples of the generated undistortion maps, computed associating at each pixel $(u, v)^\mathsf{T}$ the value $\mathrm{u}_{u,v}(d) - d$ for a given d, are reported in Fig. 16: obviously, the maps become smoother as the bin size increases.

In our experience, we found that a bin size of 4×4 pixels represents a good trade-off between computational efficiency and robustness. Actually, as mentioned before, larger bins tend to fail close to the image corners, while in other experiments

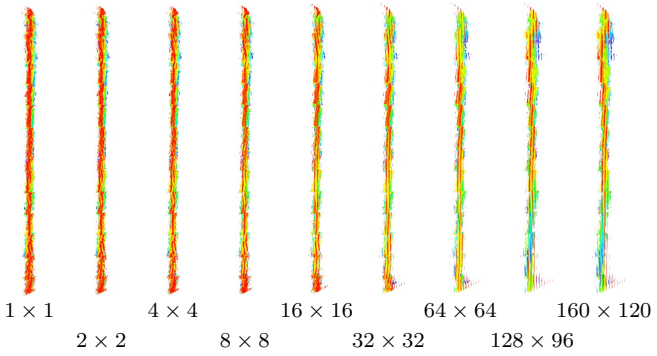

Fig. 15. Top-view of the cloud of a planar surface undistorted using maps with different bin sizes. Note that the resulting clouds are similar, especially the four on the left (points with different y-coordinates are drawn with different colors).

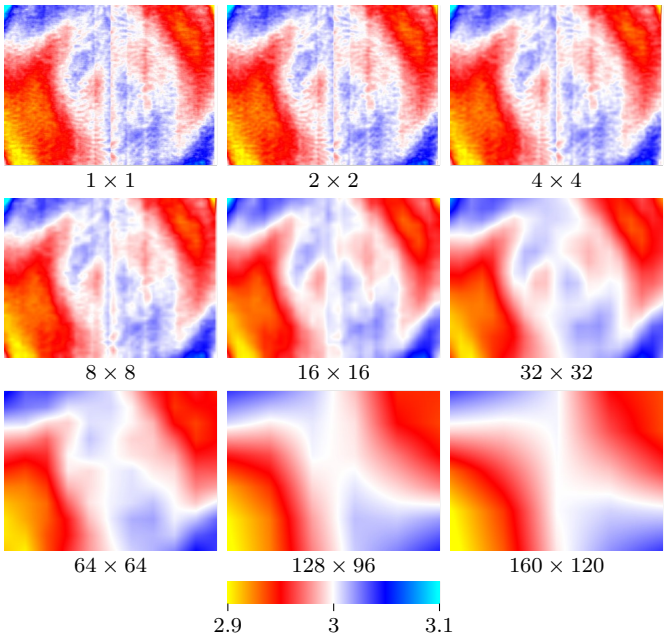

Fig. 16. Undistortion maps computed using different bin size, evaluated at a distance of 3 meters.

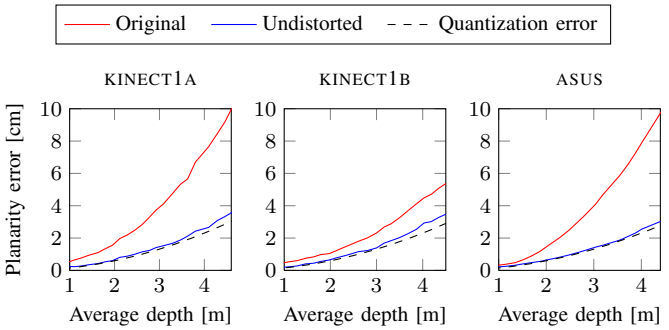

Fig. 17. Planarity error for the three tested sensors. As visible, the proposed approach is able to drastically reduce the distance of the measured points from the plane that best fits the data. The error after the undistortion phase is mainly due to quantization [13].

we noticed that smaller bins tend to perform badly with small calibration datasets because of the lack of data for some pixels.

3) Test Set Results: We finally tested our algorithm against the test sets acquired for the three SL sensors used in the experiments. Results of the planarity error evaluation are reported in Fig. 17. As expected, the proposed method permits to drastically improve the planarity of the depth data generated by the calibrated sensor. Looking at the plots one could argue: why isn't the error after the undistortion closer to zero? Actually, to the best of our knowledge, the error curve calculated after the undistortion is mainly due to the sensor noise and the quantization error. Therefore it is not possible to further reduce the error.

In Fig. 18 the estimated undistortion maps, evaluated at 4 defined depths, are shown. Looking at the scales, we can see that the magnitude of the correction is consistent with the planarity error of the original data. In Fig. 19 the results of our undistortion algorithm applied to the clouds of Fig. 1 are reported. As expected, the clouds are now planar but they are not in the right position, not even correctly oriented. These errors will be corrected in the next stage of our algorithm.

B. Global Correction Map

In this section we report some results of the evaluation of the estimated global correction map G. Recalling that G is computed using the checkerboard as the reference plane and that G and the RGB camera-depth sensor transformation are refined together, also the estimated transformation is taken into account to evaluate the results. We estimate the error of the plane that results after the global correction with respect to the plane defined by the checkerboard.

Firstly, the pose of the checkerboard with respect to the camera is estiamted using the corners extracted from the image. Then, the checkerboard plane is transformed into depth sensor coordinates. Finally, the average distance of the wall points (extracted from the cloud) to the checkerboard-defined plane is computed.

1) Global Correction Map Functions: Before evaluating the global correction map G, as we did for the undistortion map U, we analyze the sample sets generated to compute the map to evince the most appropriate function type to fit to

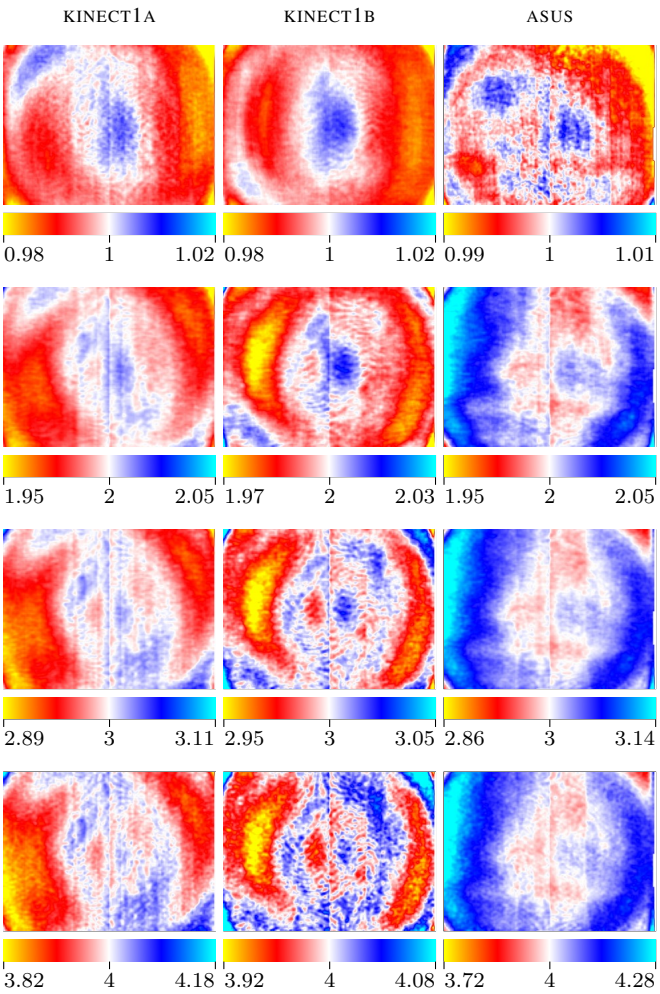

Fig. 18. Maps generated for the three sensors. For each sensor the evaluation of the respective map at 1, 2, 3, and 4 meters is reported. Note that each map has its own scale and all values are in meters.

the data. A first analysis of the error has been reported in Sect. III, in Fig. 2a. Such error was computed by evaluating the difference between the average depth of a distorted point cloud and the measurements provided by a laser distance meter. Our sample sets, instead, contain pairs $(\widehat{d}, d_{\pi_B})$, where $\widehat{d}$ is the depth value after the undistortion phase, and d_{π_B} is the expected depth, i.e. the depth of the pixel as if it were laying on checkerboard-defined plane π_B. Even in this case we tried to fit polynomial functions with different degrees: in Fig. 20 we report a comparison between quadratic and linear functions for two of the four sample sets used to generate the global correction map. From the figure, it is clear that linear functions are not suitable since they do not fit properly to the data. A further confirmation of this fact is visible in Fig. 21, where the results of the calibration process varying the maximum degree of the polynomials from 1 to 4 is reported.

Moving a cloud along the z-axis is the same as moving one of the sensors along the same axis. For this reason, to avoid problems in the optimization phase (i.e. one parameter can be increased and the other decreased of the same amount and the error evaluation does not change), for these polynomials we set the constant factor to zero.

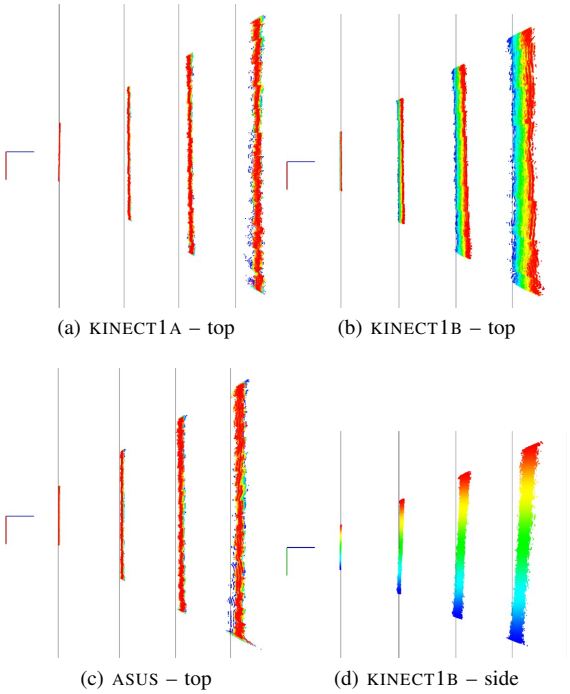

(a) KINECT1A – top (b) KINECT1B – top

(c) ASUS – top (d) KINECT1B – side

Fig. 19. Top and side views of the point clouds of Fig. 1 after the undistortion phase. Again, the gray lines show the depth measured by means of the laser meters, while points with different y-coordinates are drawn with different colors. As we can see, the clouds are now more planar than the original ones, however, they are not in the right position, not even correctly oriented.

All the calibrations reported in the following were performed treating each $g_{u,v}(\cdot)$ as a quadratic function with the constant factor set to zero.

2) Test Set Results: We finally estimated the global correction results on the acquired test sets. For each test cloud we evaluated the distances of the points of the main plane to the plane defined by the checkerboard and computed their mean. The plots in Fig. 22 show the results of such evaluation. The proposed error correction approach is working as expected: all the points are correctly translated to the right location, with respect to the checkerboard pose.

C. Testing the Whole Procedure

The following tests are meant to evaluate the results of the proposed calibration approach when dealing with real world data. To this aim, we first compare the wall average depths obtained after the calibration with the measurements given by the laser meters, then the transformations between the depth sensors and the cameras are evaluated in terms of visual results and expected values.

1) Depth Calibration: The plots in Fig. 23 show a quantitative evaluation of the depth error, i.e. the distance between the wall measured by means of the laser meters and the average depth of the points after both the undistortion and global correction phases. The measured planes are within a couple of centimeters from the real ones: these good results confirm the soundness of our choices.

We also evaluated how much the resulting plane is rotated with respect to the real one. To this aim we computed the

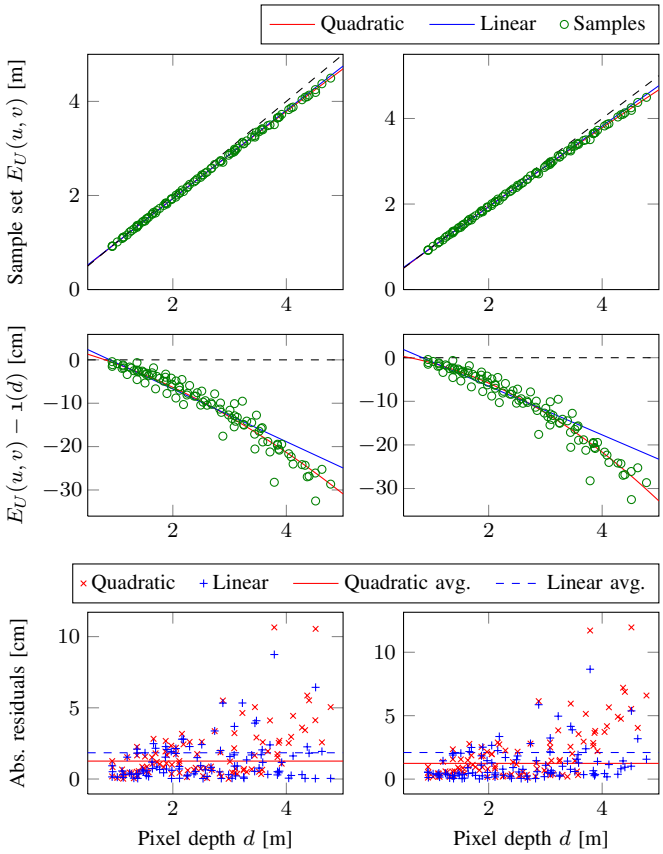

Fig. 20. Analysis of two of the sample sets used to estimate the global correction map G. The sample set is composed by pairs $(x, y) = (\widehat{d}, d_B)$, where $\widehat{d}$ is the sensor-provided depth while d_B is the depth defined by the checkerboard $\mathcal{B}$ attached on the wall. The plots at the top show the samples as well as a linear and a quadratic function that best approximate them. The plots in the middle, instead, show the difference between the samples' y and x values, i.e. $d_B - \widehat{d}$. Finally, in the bottom plots, the residuals for the two functions are reported.

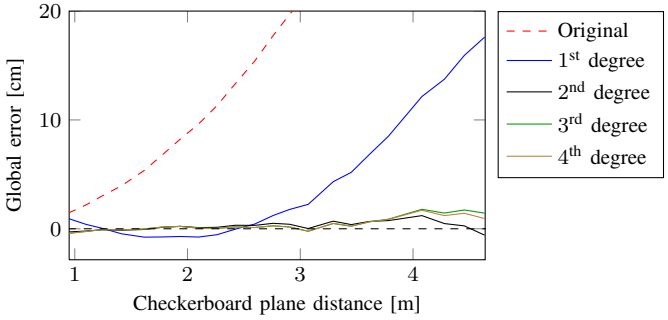

Fig. 21. Global error when varying the maximum degree of the global correction polynomials. From the plot we can see that linear functions are not able to model (and correct) the error on the average depth estimation.

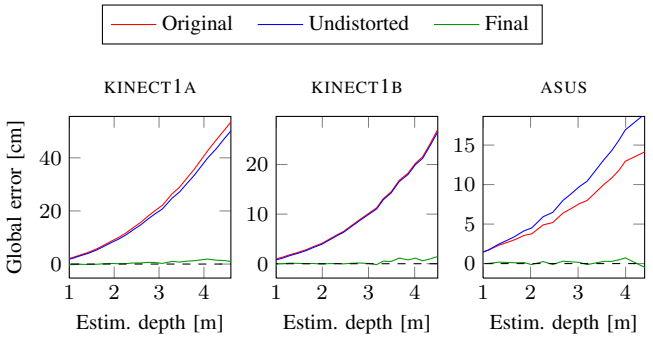

(a) Calibration using the device camera.

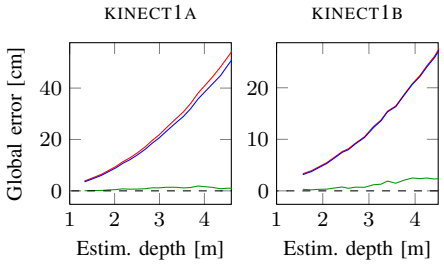

(b) Calibration using the external, high resolution camera.

Fig. 22. Global error for the three tested SL depth sensors. We report the error for the original point cloud (Original), the error after the undistortion step (Undistorted), and the error after both the undistortion and the global error correction steps (Final). To further assess the validity of the proposed approach, in (a) we calibrated the depth sensor using the device camera, while in (b) we used the external high resolution camera (see Fig. 11a).

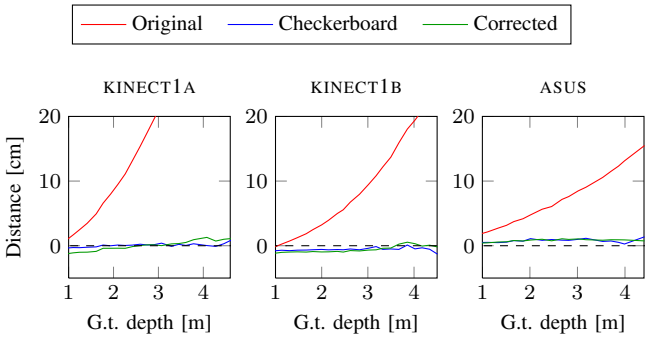

(a) Results of the calibration using the device camera.

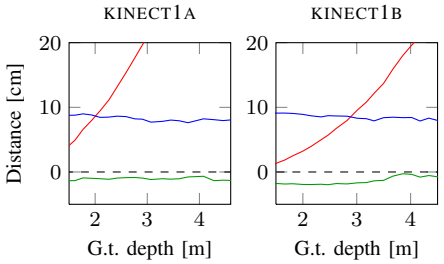

(b) Results of the calibration using the external, high resolution camera.

Fig. 23. Distance between the real wall depth and the one estimated with the calibration procedure. The error is computed for the original point clouds (Original) and for the clouds after both the undistortion and the global error correction (Corrected). Moreover, the distance of the wall from the color sensor estimated using the checkerboard (Checkerboard) is reported. The error is computed using the device camera (a) as the reference camera as well as using the external high resolution one (b). Note that there is a fixed offset of less than 1 cm between the laser meters and the two Kinects (KINECT1A and KINECT1B are closer to the wall) and about 9 cm between the laser meters and the high resolution camera (the camera is farther).

angle between the normal of the plane fitted to the corrected data, and the x- and y-axis of the wall plane, i.e. $(1,0,0)^\mathsf{T}$ and $(0,1,0)^\mathsf{T}$ respectively. Let $\mathbf{n}$ be the fitted-plane normal and let $\mathbf{a}$ be the axis with respect to which the error is computed, the rotation error, e_{rot}, is

$$e_{\text{rot}} = \arccos(\mathbf{n}^\mathsf{T}\mathbf{a}) - \frac{\pi}{2} \ .$$

Results of the error computation for sensor KINECT1B are reported in Fig. 24. The figure shows that the rotation about the x-axis is completely corrected. For what concerns the rotation about the y-axis, instead, the results are worse. The reason for this fact is likely to be the error on the real depth estimation. In fact, a difference of about 2 mm in the depth measures (note that this is the nominal error of the two laser meters) leads to a rotation of about 0.5°.

A further confirmation that the proposed approach works properly, is shown in Fig. 25. The pictures report the clouds of Fig. 1 after both the undistortion phase (see Fig. 19) and the global error correction. As expected, all the clouds are now both planar and located correctly.

2) Camera-Depth Sensor Transform: Even if the camera-depth sensor transform ${}^{\mathcal{D}}_{\mathcal{C}}\mathbf{T}$ estimated during the optimization phase is a sort of "side effect" of the depth calibration, a good transformation can be seen as a proof of the validity of the proposed approach. In Tab. I the transformations that resulted from the calibration of the sensors using their device cameras, are reported. Moreover, to give a comparison metric, also the factory-provided transformation is reported. The values

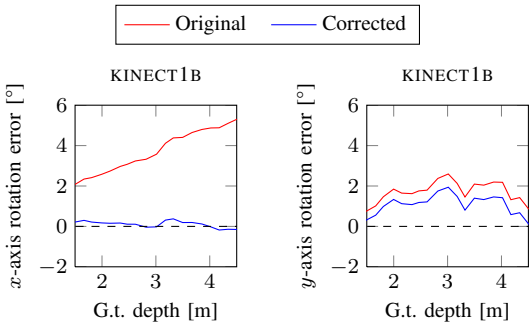

Fig. 24. Rotation error for KINECT1B sensor. Both the rotation about the x- and y-axis are compared. The error is the angle estimated between the normal of the plane fitted to the corrected data and the theoretical wall plane x-axis and y-axis.

obtained are similar to those obtainable with other state-of-the-art calibration tools for RGB-D devices [18], [34], [19], [15].

Some qualitative results of this camera-depth sensor calibration are shown in Fig. 26. Looking at the images, we can see that the data fusion obtained with the transformation computed during the calibration process are definitely better than the ones obtained with the default transformation.

TABLE I

CAMERA-DEPTH SENSOR TRANSFORM ESTIMATED FOR THE THREE TESTED SENSORS WHEN CALIBRATED USING THE DEVICE CAMERA. BOTH THE TRANSLATION $\mathbf{t} = (t_x, t_y, t_z)^\mathsf{T}$ AND THE ROTATION, REPRESENTED AS A QUATERNION $\mathbf{q} = (q_w, q_x, q_y, q_z)^\mathsf{T}$, ARE REPORTED. THE FACTORY LINE CONTAINS THE FACTORY-PROVIDED CALIBRATION PARAMETERS FOR THE THREE DEVICES.

	t_x [m]	t_y [m]	t_z [m]	q_x	q_y	q_z	q_w
FACTORY	0.025	0	0	0	0	0	1
KINECT1A	0.0237	0.0044	-0.0063	0.0034	0.0060	-0.0017	0.9999
KINECT1B	0.0276	0.0024	-0.0036	0.0025	0.0007	-0.0010	0.9999
ASUS	0.0294	-0.0040	-0.0011	0.0048	0.0059	-0.0004	0.9999

TABLE II

COMPARISON OF THE CALIBRATION ACCURACY FOR A KINECT 1 SENSOR.

	$\mu(\varepsilon_3)$ [m]	$\sigma(\varepsilon_3)$ [m]	$\mu(\varepsilon_2)$ [pixels]	$\sigma(\varepsilon_2)$ [pixels]	$\mu(\measuredangle_x)$ [°]	$\mu(\measuredangle_y)$ [°]	$\mu(\measuredangle_z)$ [°]
ORIGINAL	0.129	0.059	5.836	0.955	1.596	1.541	1.796
HERRERA *et al.* [18]	0.028	0.017	2.388	0.800	0.814	1.258	1.404
STARANOWICZ *et al.* [19]	0.172	0.056	4.246	1.254	1.364	1.186	1.842
OUR METHOD	**0.011**	**0.004**	**1.901**	**0.717**	**0.691**	**0.617**	**0.930**

TABLE III

COMPARISON OF THE CALIBRATION ACCURACY FOR A KINECT 2 SENSOR.

	$\mu(\varepsilon_3)$ [m]	$\sigma(\varepsilon_3)$ [m]	$\mu(\varepsilon_2)$ [pixels]	$\sigma(\varepsilon_2)$ [pixels]	$\mu(\measuredangle_x)$ [°]	$\mu(\measuredangle_y)$ [°]	$\mu(\measuredangle_z)$ [°]
ORIGINAL	0.096	**0.008**	20.913	**2.310**	2.084	2.308	2.424
OUR METHOD	**0.075**	0.011	**9.272**	3.473	**1.685**	**1.339**	**1.896**

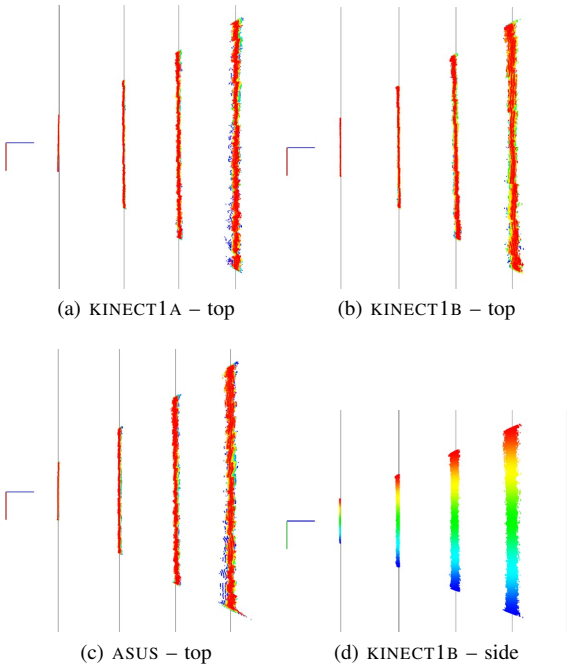

(a) KINECT1A – top (b) KINECT1B – top

(c) ASUS – top (d) KINECT1B – side

Fig. 25. Top and side views of the point clouds of Fig. 1 after the undistortion phase (Fig. 19) and the global correction phase. Again, the gray lines show the depth measured by means of the laser meters. As we can see, every cloud is now planar as well as in the right position and correctly oriented.

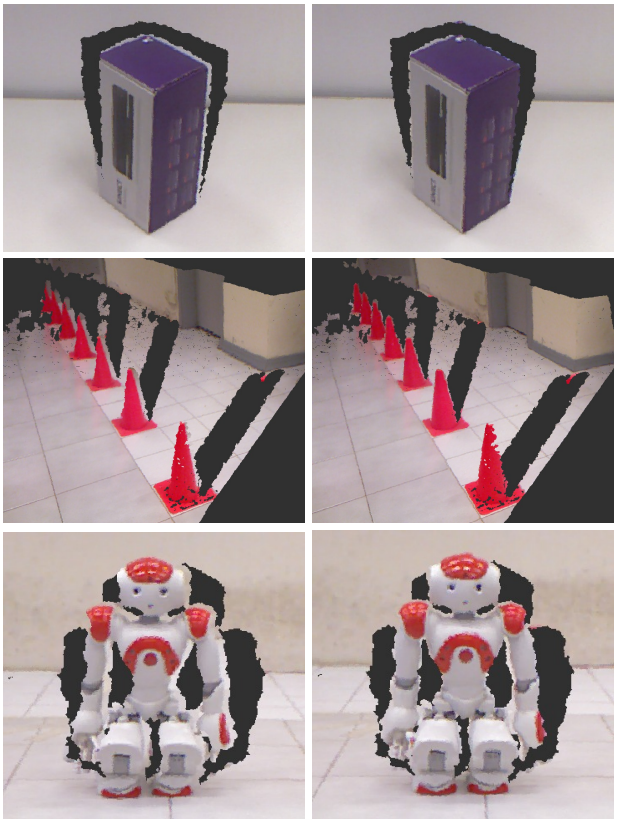

(a) Default transformation (b) Computed transformation

Fig. 26. Colored point clouds resulting from the fusion of depth and RGB data. For each cloud, the points are firstly transformed into camera coordinates, then projected into the image and colored with the RGB color of the obtained pixel. As we can see, the results obtained with the newly computed transformation (b) are definitely better than those obtained with the default one (a).

D. Performance Comparison

We tested the calibration accuracy of our system against two state-of-the-art calibration methods, the one from Herrera *et al.* [18] and the one from Staranowicz *et al.* [19], using the original implementations provided by the authors. Thanks to its robustness and accuracy, [18] is often used as a benchmark-

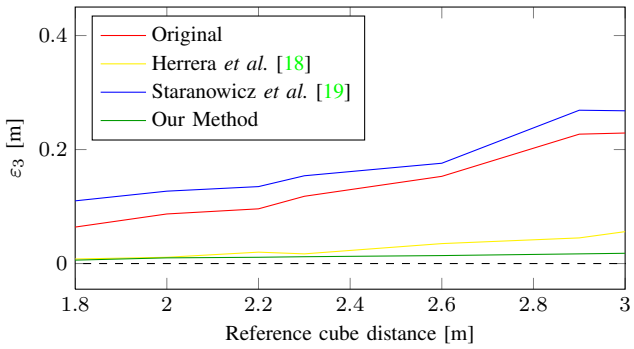

Fig. 27. Accuracy in the reference cube localization for increasing depths.

ing method to calibrate RGB-D pairs; [19] is a recent method that, compared to [18] and our method, employs a novel, alternative calibration procedure based on a spherical pattern. We used two different types of RGB-D sensors: a Kinect 1 and a Kinect 2. For each method, we acquired large training sets, framing the calibration pattern (a wall for [18] and our method, a basket ball for [19]) from several positions. In order to compare the calibration accuracy, we collected a test set framing a big reference hollow cube with large checkerboards attached to each visible side (Fig. 11b). Using this 3D pattern, it is possible to accurately compute the planes equations of the three sides using the corners extracted from the image: we use these planes, their intersection point $\mathbf{x}_c$ and its projection in the image plane $\mathrm{proj}_{\mathcal{C}}(\mathbf{x}_c)$ as ground truth data. For each tested method, we estimated the plane equations by fitting the three planes to the (corrected) point clouds, computing also their intersection point $\mathbf{x}'_c$ as well as its projection in the image plane $\mathrm{proj}_{\mathcal{C}}(\mathbf{x}'_c)$. In our performance comparison, we report the average $\mu(\cdot)$ and standard deviation $\sigma(\cdot)$ of both the errors $\varepsilon_3 = \|\mathbf{x}_c - \mathbf{x}'_c\|$ and $\varepsilon_2 = \|\mathrm{proj}_{\mathcal{C}}(\mathbf{x}_c) - \mathrm{proj}_{\mathcal{C}}(\mathbf{x}'_c)\|$, computed for all the images included in the test set. We also report the average of the angular deviations $\measuredangle_x, \measuredangle_y, \measuredangle_z$ between the ground truth planes and the estimated planes. For a perfect calibration, these errors should be obviously 0.

Tab. II shows the performance comparison results of the three tested method for a Kinect 1 sensor; as baseline, we also report the results obtained using the original, factory-calibrated data. Fig. 27 shows error ε_3 for increasing depths from the reference cube. Our method clearly provide the better calibration accuracy. The good results obtained by the method from Herrera *et al.* confirm that the choice of a planar calibration pattern enables to obtain superior calibration results. The poor results obtained by the method from [19] are mainly due to the inaccuracies of the calibration pattern detector: we collected several training sets, framing a basket ball (as suggested by the authors) from a large number of different positions, but the provided ball detector often failed to provide an accurate localization.

We also compared the calibration accuracy of our method against the factory calibration in the case of a Kinect 2 sensor (Tab. III). The method from Herrera *et al.* is here not applicable since it operates directly on the disparity map provided by a SL sensor: clearly a ToF sensor does not provide such map. Also in this case, our method outperforms the factory calibration, enabling to obtain better data also in the case of a ToF sensor. Surprisingly the Kinect 2 performed significantly worse than the older Kinect 1: this result can be explained by the fact that ToF sensors work very poorly when dealing with black surfaces as the black squares of the checkerboards attached to the reference cube.

E. Visual Odometry Use Case

As a further validation of our method, we present an experiment performed using a real robot running an RGB-D visual odometry system. Here we show how the accuracy in the ego-motion estimation and the 3D reconstruction can highly benefit from using RGB-D data calibrated with our method. We employed two different RGB-D visual odometry systems: the popular and robust DVO (Dense Visual Odometry) [35] and a very simple custom-built system based on dense optical flow, in the following OFVO (Optical Flow Visual Odometry). DVO registers two consecutive RGB-D frames by minimizing the photometric error between corresponding pixels while OFVO firstly computes the dense optical flow between a reference and the current RGB image and, after generating the point cloud of the reference frame using the depth image, it estimates the relative rigid motion motion by solving a Perspective-n-Point problem[4]. These methods strongly rely on both the intrinsic and extrinsic calibration of the RGB-D pair, so they represent a perfect benchmark for our method which provides the complete calibration for such sensors.

We moved a mobile robot (a MobileRobots Pioneer 3-AT) equipped with a Microsoft Kinect 1 along a known trajectory, then we estimated the motion[5] using both DVO and OFVO on the original RGB-D data and on the same data correct with our method. Table Tab. IV shows the root mean square error (RMSE) of the estimated motion for each experiment: in both cases the accuracy improvement is remarkable. Fig. 28(a) shows a top view of the estimated and ground truth trajectories, along with the generated point clouds (Fig. 28(b), (c)). For both methods, the trajectory estimated using the corrected data is clearly closer to the ground truth compared to the one estimated using the original data: in the first case, most of the misalignment is mainly due to a drift accumulated in the first turn, where the robot acquired a sequence of frames with very few visual features. The trajectory estimated using the original data tends to diverge also due to a sort of "scale drift" effect that supports our choice to introduce the global component in our error model. Similarly, the quality and the precision of the reconstructed point cloud highly benefits from using the corrected data.

F. Runtime Performance

The proposed algorithm has been implemented in C++ as a ROS [6] package[6], using the OpenCV [36] library for image processing, the Point Cloud Library (PCL) [30] for 3D data

[4]To solve this problem, we used the OpenCV function *solvePnP*.

[5]We estimated the *full* 3D motion, without using any planar motion assumption or wheel odometry prior.

[6]http://iaslab-unipd.github.io/rgbd_calibration

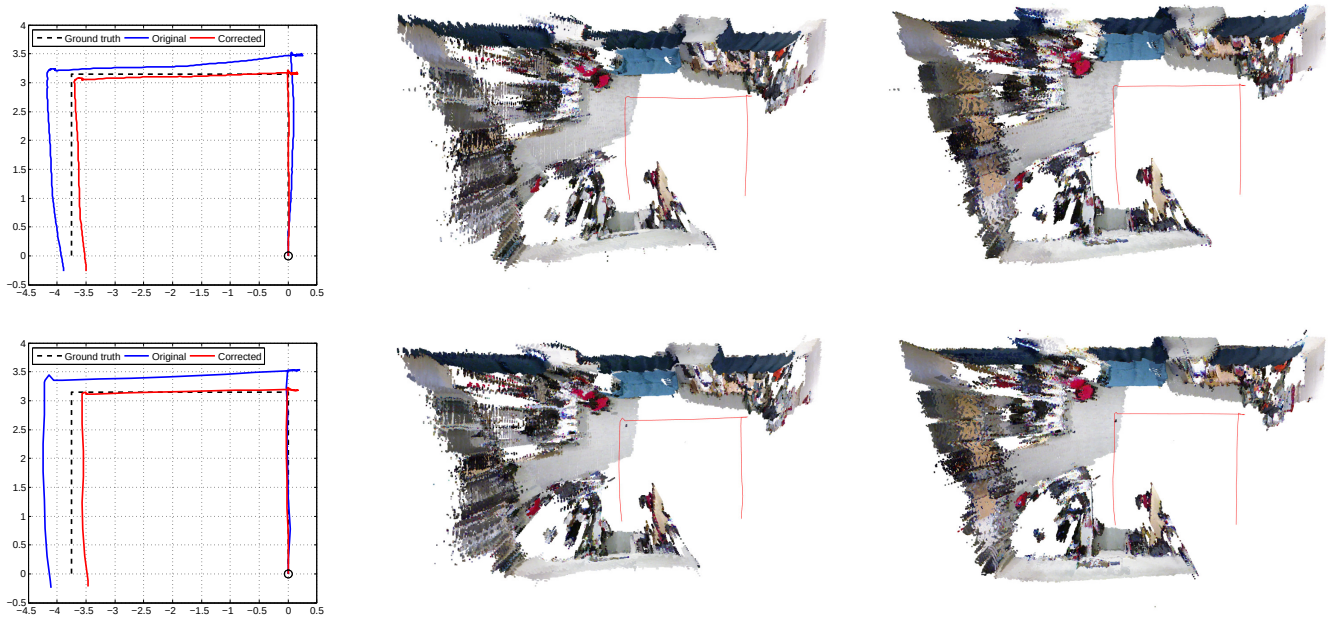

(a) Top view of the trajectories. The start location is surrounded with a circle.

(b) Point cloud obtained using the original data.

(c) Point cloud obtained using the corrected data.

Fig. 28. Qualitative results of the visual odometry experiments: (first row) DVO [35] visual odometry systems; (second row) OFVO visual odometry systems.

TABLE IV
RMSE OF THE DRIFT IN METERS FOR THE WHOLE PATH.

	Original data [m]	Corrected data [m]
DVO [35]	0.3371	**0.1589**
OFVO	0.3530	**0.2501**

processing, and the Ceres Solver library [37] to solve the optimization problems. To perform the calibration, the user is asked to capture a training set using a tool provided in the package, by moving the sensor in front of a wall with a checkerboard attached on it. Collecting a typical dataset of a hundred images takes no more than 10 minutes.

We tested both the calibration and the correction stages in terms of the execution time. The whole calibration process takes about 45 minutes on a consumer laptop[7] for 640×480 depth images and 5 minutes by downsampling the depth images to a resolution of 320×240 pixels: it is worth to mention that in the last case, the calibration accuracy is not significantly reduced.
The calibration is an operation that is performed once and so the execution time is not critical. On the other side, the execution time of the correction stage node is critical, since the data generated by the RGB-D sensor, typically with a frequency of 30 Hz, should be processed in real-time. We tested the performance of 3 different implementations of the correction algorithm:

- a standard CPU implementation;
- a parallel CPU implementation exploiting the OpenMP directives;
- a parallel GPU implementation using CUDA.

[7]CPU: Intel Core i7-4700MQ, RAM: 16GB, SSD, GPU: NVidia GTX 750M.

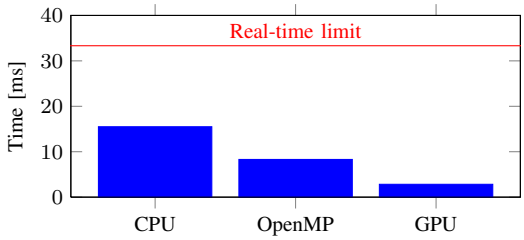

Fig. 29. Time comparison between the 3 different implementations of the correction node. The red line indicates the time limit to elaborate point clouds at 30Hz.

The results are reported in Fig. 29: note that these execution times include the time to generate the (corrected) point cloud from the depth image. Clearly the GPU implementation outperforms the CPU ones, but dedicated hardware is needed. Most of the time (about 95%) spent by the GPU implementation is dedicated to the copy of the data to and from the GPU memory.
Anyway, all the implementations are able to correct the data in real-time.

VIII. CONCLUSIONS

In this paper we presented a novel method to calibrate a general RGB-D sensor. The proposed calibration procedure only requires the user to collect data in a minimally structured environment, providing in output both the intrinsic and extrinsic parameters of the sensor. We proposed to model the

depth sensor error by means of two different components, a distortion error and a global, systematic error. The distortion error is modeled using a per-pixel parametric undistortion map, estimated in the first stage of the algorithm. The depth systematic error along with the camera-depth sensor alignment are estimated in the second stage of the algorithm, inside a robust optimization framework. We reported a comprehensive set of tests that support the introduced model. We finally presented exhaustive experiments performed using several sensors, showing that our approach provides highly accurate results, outperforming other state-of-the-art methods.
Finally, the proposed system has been released as open source package within the robotics framework ROS.

IX. Acknowledgments

The research work is partially supported by: the European Commission under 601116-ECHORD++ (FlexSight experiment) and the University of Padua under the project DVL-SLAM.

References

[1] F. Endres, J. Hess, J. Sturm, D. Cremers, and W. Burgard, "3d mapping with an RGB-D camera," *IEEE Transactions on Robotics (T-RO)*, vol. 30, no. 1, pp. 177–187, 2013.

[2] M. Labb and F. Michaud, "Online global loop closure detection for large-scale multi-session graph-based SLAM," in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2014, pp. 2661–2666.

[3] M. Munaro and E. Menegatti, "Fast RGB-D People Tracking for Service Robots," *Autonomous Robots*, 2014.

[4] C. C. and H. Christensen, "RGB-D object tracking: A particle filter approach on gpu," in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 1084–1091.

[5] J. Tang, S. Miller, A. Singh, and P. Abbeel, "A textured object recognition pipeline for color and depth image data," in *Proc. of the IEEE International Conference on Robotics and Automation*, 2012.

[6] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," in *ICRA Workshop on Open Source Software*, 2009.

[7] C. Mei and P. Rives, "Calibration between a central catadioptric camera and a laser range finder for robotic applications," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2006, pp. 532–537.

[8] D. Scaramuzza, A. Harati, and R. Siegwart, "Extrinsic Self Calibration of a Camera and a 3D Laser Range Finder from Natural Scenes," in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2007, pp. 4164–4169.

[9] Y. M. Kim, D. Chan, C. Theobalt, and S. Thrun, "Design and calibration of a multi-view ToF sensor fusion system," in *Proc. of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2008, pp. 1–7.

[10] J. Jung, Y. Jeong, J. , H. Ha, D. J. Kim, and I. Kweon, "A novel 2.5D pattern for extrinsic calibration of ToF and camera fusion system," in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 3290–3296.

[11] N. Burrus. (2011) Kinect calibration. [Online]. Available: http://nicolas.burrus.name/index.php/Research/KinectCalibration

[12] K. Konolige and P. Mihelich. (2011) Technical description of kinect calibration. [Online]. Available: http://www.ros.org/wiki/kinect_calibration/technical

[13] J. Smisek, M. Jancosek, and T. Pajdla, "3D with Kinect," in *Proc. of IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2011, pp. 1154–1160.

[14] K. Khoshelham and S. O. Elberink, "Accuracy and resolution of kinect depth data for indoor mapping applications," *Sensors*, vol. 12, no. 2, pp. 1437–1454, 2012.

[15] C. Zhang and Z. Zhang, "Calibration between depth and color sensors for commodity depth cameras," in *Proc. of IEEE International Conference on Multimedia and Expo (ICME)*, 2011, pp. 1–6.

[16] I. V. Mikhelson, P. G. Lee, A. V. Sahakian, Y. Wu, and A. K. Katsaggelos, "Automatic, fast, online calibration between depth and color cameras," *Journal of Visual Communication and Image Representation*, vol. 25, 2014.

[17] D. Herrera C., J. Kannala, and J. Heikkilä, "Accurate and practical calibration of a depth and color camera pair," in *Computer Analysis of Images and Patterns*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2011, vol. 6855, pp. 437–445.

[18] ——, "Joint depth and color camera calibration with distortion correction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 10, pp. 2058–2064, 2012.

[19] A. N. Staranowicz, G. R. Brown, F. Morbidi, and G. L. Mariottini, "Practical and accurate calibration of RGB-D cameras using spheres," *Computer Vision and Image Understanding*, vol. 137, pp. 102–114, 2015.

[20] W. Xiang, C. Conly, C. D. McMurrough, and V. Athitsos, "A review and quantitative comparison of methods for kinect calibration," in *Proc. of the 2nd International Workshop on Sensor-based Activity Recognition and Interaction (WOAR)*, 2015.

[21] C. Raposo, J. P. Barreto, and U. Nunes, "Fast and accurate calibration of a kinect sensor," in *Proc. of International Conference on 3D Vision (3DV)*, June 2013, pp. 342–349.

[22] A. Canessa, M. Chessa, A. Gibaldi, S. P. Sabatini, and F. Solari, "Calibrated depth and color cameras for accurate 3D interaction in a stereoscopic augmented reality environment," *Journal of Visual Communication and Image Representation*, vol. 25, no. 1, pp. 227–237, 2014.

[23] A. Teichman, S. Miller, and S. Thrun, "Unsupervised intrinsic calibration of depth sensors via SLAM," in *Proc. of Robotics: Science and Systems (RSS)*, Berlin, Germany, June 2013.

[24] D. Fiedler and H. Müller, "Impact of thermal and environmental conditions on the kinect sensor," in *Advances in Depth Image Analysis and Applications*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2013, vol. 7854, pp. 21–31.

[25] M. Di Cicco, L. Iocchi, and G. Grisetti, "Non-parametric calibration for depth sensors," *Robotics and Autonomous Systems*, vol. 74, pp. 309–317, 2015.

[26] F. Basso, A. Pretto, and E. Menegatti, "Unsupervised intrinsic and extrinsic calibration of a camera-depth sensor couple," in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 6244–6249.

[27] E. Lachat, H. Macher, M.-A. Mittet, T. Landes, and P. Grussenmeyer, "First Experiences with Kinect v2 Sensor for Close Range 3d Modelling," *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pp. 93–100, Feb. 2015.

[28] D. Pagliari and L. Pinto, "Calibration of kinect for xbox one and comparison between the two generations of microsoft sensors," *Sensors*, vol. 15, no. 11, pp. 27 569–27 589, 2015.

[29] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, June 1981.

[30] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011, pp. 1–4.

[31] A. Kaehler and G. Bradski, *Learning OpenCV, 2nd Edition*. O'Reilly Media, 2014.

[32] R. Unnikrishnan and M. Hebert, "Fast extrinsic calibration of a laser rangefinder to a camera," Carnegie Mellon University, Tech. Rep., 2005.

[33] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment – a modern synthesis," in *Vision Algorithms: Theory and Practice*, ser. Lecture Notes in Computer Science, B. Triggs, A. Zisserman, and R. Szeliski, Eds. Springer Berlin Heidelberg, 2000, vol. 1883, pp. 298–372.

[34] A. Staranowicz, G. R. Brown, F. Morbidi, and G. L. Mariottini, "Easy-to-use and accurate calibration of RGB-D cameras from spheres," in *Image and Video Technology*, ser. Lecture Notes in Computer Science, R. Klette, M. Rivera, and S. Satoh, Eds. Springer Berlin Heidelberg, 2014, vol. 8333, pp. 265–278.

[35] C. Kerl, J. Sturm, and D. Cremers, "Robust odometry estimation for rgb-d cameras," in *Proc. of the IEEE International Conference on Robotics and Automation*, 2013.

[36] G. Bradski, "The OpenCV library," *Dr. Dobb's Journal of Software Tools*, 2000.

[37] S. Agarwal, K. Mierle, *et al.*, "Ceres solver," http://ceres-solver.org.